Working Paper Series

Risk, the College Premium, and Aggregate Human Capital Investment

WP 13-02R

Kartik Athreya
Federal Reserve Bank of Richmond

Janice Eberly
Northwestern University and NBER

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ■ BALTIMORE ■ CHARLOTTE

Risk, the College Premium, and Aggregate Human Capital Investment *

Kartik Athreya[†] Janice Eberly[‡§]

November 13, 2018

Working Paper No. 13-02R

Abstract

Given the large and rising return to education, the response of human capital investment has been described as comparatively "anemic". This paper shows that dispersion in the likelihood of college noncompletion, and to a lesser extent risks to earnings over the life cycle, may strongly limit the response of aggregate human capital investment to further increases in the U.S. college earnings premium. Under educational conditions reflecting those prevailing currently in the U.S., our results suggest that the continuation of long-standing trends in the skill-bias of technical change (Goldin and Katz, 2008) can be expected primarily to increase earnings and income inequality, rather than college attainment.

Keywords: College Enrollment, College Non-Completion, College Premium, Risk. **JEL Codes**: I21, I24, J24

1 Introduction

This paper demonstrates that dispersion in the likelihood of college noncompletion, and to a lesser extent risks to earnings over the life cycle, may strongly limit the response of aggregate human capital investment to further increases in the U.S. college earnings premium. Under educational conditions reflecting those prevailing currently in the U.S., our results suggest that the continuation of long-standing trends in the skill bias of technical change (Goldin and Katz, 2008) can be expected primarily to increase earnings and income inequality, rather than college attainment.

The starting point for our analysis is the observation that the earnings differential between college and high school graduates is persistently large and has increased even further in recent years. At current levels, for example, a college graduate earns nearly twice as much over a lifetime (in present value) as a high school graduate (e.g., Restuccia and Urrutia, 2004; Hendricks and Leukhina, 2012). However, as this "college

*We thank the editor, Richard Rogerson, and two anonymous referees for helpful comments on the manuscript, as well as Felicia Ionescu, John Bailey Jones, John Knowles, Bruce Sacerdote, and Sarah Turner for detailed comments. We also thank Gonzalo Castex, Yongsung Chang, Stacey Chen, Doug Diamond, Caroline Hoxby, Fabian Lange, Debbie Lucas, Ariell Reshef, Cecilia Rouse, Felipe Schwartzman, seminar participants at Colgate, Stanford, Virginia, and the Kansas City, New York, and Richmond Feds for their feedback. Brian Gaines, Jonathan Tompkins, and especially Roisin McCord provided exceptional research assistance.

[†]Federal Reserve Bank of Richmond.

[‡]Kellogg School of Management, Northwestern University, and NBER

premium" has grown, once rapid growth in college attainment has slowed (Altonji, Bharadwaj, and Lange, 2008b; Goldin and Katz, 2008). Notably, Castro and Coen-Pirani (2015) show that cohort-level college attainment rates have remained essentially constant for two decades and argue that reductions in a measure of *ability* account for roughly half of the stagnation in the 1972 birth cohort relative to its 1948 counterpart (with changes in the cost of college accounting for the rest). These findings are suggestive of a bottleneck in the supply of college-completed individuals stemming from the non-college-bound population facing lower returns to collegiate attainment.

Our paper argues that selection by college preparedness (as measured by the noncompletion probability) plays an important role in lowering these returns and hence in limiting further increases in attainment in response to increases in the college premium.

While preparedness for higher education may be culpable in these trends, it can be abstract and difficult to measure. The literature typically proxies preparedness with measured test scores. However, one aspect of preparedness is, almost by definition, plainly observable: whether an individual completes a course of study. In recent cohorts, collegiate noncompletion has been high, with roughly 50 percent of all public four-year entrants failing to secure any degree within eight years of entry (Bound, Loevenheim, and Turner, 2009). Even conditional on completion, moreover, U.S. households face significant idiosyncratic risk that appears essentially uninsurable (see, e.g., Heathcote, Storesletten, Violante, 2009). In the context of human capital investments such as college, this includes adverse labor market events that can greatly diminish, if not fully eliminate, the ex-post earnings advantage realized by those who enroll in and finish college.

These disparate forces all lead, plausibly, to investment in college being both a *low expected return* one and a risky one as well– arguably for a large fraction of high school completers. Our interest lies in answering two questions: First, how is college enrollment influenced by the earnings premium in an environment with realistic uncertainty around both completion and then subsequent earnings? Second, given that noncompletion rates may change with enrollment, what are the implications for aggregate U.S. college attainment? That is, can further increases in the college premium be expected to bring forth substantially more additional college-educated individuals?

To answer these questions, we augment a standard model of consumption and savings to allow for college investment. The structure is consistent with recent work (Stange, 2012; Ionescu and Chatterjee, 2012) and models college investment as potentially low return (students may enroll but not complete, and payoffs conditional on completion are partially random), lumpy (college returns are disproportionate upon completion), and irreversible (investment in college cannot be decumulated to finance consumption). We parameterize the model to match observed college enrollment, including both two- and four-year colleges, and completion rates under the currently prevailing college premium. We then use the model to measure the responsiveness of college enrollment and completion to changes in the college premium, holding all else fixed.

We have two central findings. First, college *enrollment* is not likely to be sensitive to small changes, especially increases, in the college premium from current levels. Second, college *attainment* will remain insensitive to even large increases in the college premium from current levels. These findings flow from an intuitive and plausible mechanism arising from the possibility of college noncompletion. Simply put, the effect of the college premium on enrollment is dampened in expectation by the possibility of noncompletion. Those students likely to complete college already enroll at high rates and would continue to do so even if the college premium rose or fell at the margin. More importantly, those potential students who do not already enroll tend to have a high probability of not completing college. An increase in the college premium creates little additional incentive to attend college for these students, since they only receive the college premium if they graduate. Hence, the lower are completion odds, the less effective is an increase in the college premium at generating greater enrollment. These effects are consistent with the empirical observations that college enrollments are currently at a historic high (at over 70 percent), while college noncompletion rates exceed 50 percent (see, e.g., Bound, Loevenheim, and Turner, 2009). In sum, our model implies strongly that the US has largely "run out of college-ready youth," which has led the relative supply curve of skilled labor to effectively be quite steep.

Our findings are relevant in part because they have immediate implications for the larger policy discussions on the distributional implications of the payoffs to human capital and the prospects for longer-run growth. Our analysis suggests that absent major changes in college preparedness, increases in college wages will, all else equal, translate into increases in overall earnings inequality in the long run. Thus, the continuation of long-running trends in skill-biased technological change (such as those arising from the falling price of computing power (Autor, Levy, and Murnane, 2003)) will likely lead to more rapid growth in income inequality than it has in the past. Recent work, including that of Goldin and Katz (2008, pp. 46-57), has stressed this possibility. As for the implications of our work for longer-run growth, Fernald and Jones (2014) identify the plateauing of educational attainment over the past several decades as important for the slowdown in future growth prospects. Our model suggests that going forward, even strong private incentives to accumulate postsecondary education, as expressed through the college premium, should not be expected to bring forth large increases in human capital, at least not through formal college education.

College attainment, which translates into human capital accumulation, is also weakly related to the college premium for similar reasons. Additional college enrollment is likely to arise, if at all, from groups that currently enroll at relatively low rates. But since such groups *complete* college at relatively low rates, the effect of the college premium on their attainment will remain muted. And because the enrollment of groups that already enroll at high rates will not change, neither can their attainment. Thus, even if the premium changes by enough to induce poorly prepared individuals to enroll at higher rates than currently observed, changes in attainment will necessarily be limited.

Our findings emerge from the interaction of several forces. Students may be well- or poorly prepared for college, face uncertain income thereafter, and also have different levels of wealth. Low odds of completion and subsequent rate-of-return risk loom largest for the poor and most poorly prepared: These groups face the possibility of being substantially worse off from the time and money spent investing in college than if they had never enrolled. Moreover, any risks inherent in college investment are further amplified by any leverage used by low-wealth households to finance college, given the presence of incomplete insurance against either noncompletion or labor earnings shocks. By contrast, while noncompletion and risky rate of return are possibilities also present for richer individuals, they can more easily self-insure so negative shocks have less serious consequences.

If the current college premium is a good measure of the expected payoff to the marginal decision maker, the absence of a substantial increase in both enrollment and the stock of college-educated labor over time seems puzzling. But as Carneiro, Heckman, and Vytlacil (2011) show, the average return to enrollment is not, in fact, the payoff to the marginal enrollee. Even if students expect to receive the currently observed premium *upon completion*, only those who successfully complete college will actually earn this premium.

However, as noted at the outset, college completion is far from universal and appears closely tied to preparation and household background (e.g., Bowen et al. 2009) As for the former, it appears that the proportion of enrollees with a low likelihood of completion has grown. Bound et al., (2009), find, for example, that for the 1992 high school graduating class in the NELS88 data set, 44 percent of those in the lowest quartile in the mathematics test enrolled in college (see their Figure 2). Yet only 11.4 percent of the 1992 cohort completed a bachelor's degree. By contrast, less than 20 years earlier, in 1976, the NLS72 data show that while the enrollment rate for such poorly prepared students was much lower, at only 21.7 percent, this group completed college at a much higher rate of 25.8 percent.[1] By various measures, the overall noncompletion rate at four-year institutions is close to 50 percent.[2]

The features we incorporate into our model of college investment are informed by two related bodies of work. First, there is abundant evidence of noncompletion and, as measured by the probability that a student will fail to complete college, in a manner not perfectly anticipated at the outset (e.g., Manski and Wise, 1983; Manski, 1989; Stange, 2012; or Arcidiacono et al., 2012).[3] The overall postsecondary noncompletion rate is currently approximately 50 percent (Singh, 2010; Bowen et al., 2009; Bound et al., 2009, Hendricks and Leukhina, 2011). The odds of completion are especially relevant because any uncertainty over eventual completion is often not quickly resolved: At four-year institutions, at the median, resolution takes between two and three years of forgone earnings and the explicit cost of tuition (Bowen et al., 2009, Ozdagli and Trachter, 2011). In addition to taking time, the return to the partial completion of college appears low (i.e., attending but not obtaining a four-year diploma); early documentation includes Layard and Psacharapoulous (1974), and more recently Hungerford and Solon (1987) and Kane and Rouse (1995). The risk of not completing school is magnified by the lumpiness of college investment along with the poor returns to noncompleters.

A second line of research stresses the presence, *conditional on college completion*, of "rate-of-return risk." Early analysis of Carneiro, Hansen, and Heckman (2003), for example, estimates that a substantial portion of the present value of earnings is unpredictable at the time individuals enroll in college and, moreover, that risk implies that the ex-post returns to college will be substantially negative for some.[4] By many accounts, households appear to face persistent uninsurable earnings shocks (e.g., Hryshko, 2010; Storesletten, Telmer, and Yaron, 2004). It is therefore entirely possible for relatively young college graduates to receive earnings shocks that immediately and substantially lower the expected present value of remaining lifetime income. The persistence of these shocks also makes them inherently difficult to self-insure. Few papers feature both noncompletion *followed by* rate-of-return risk. The exceptions are recent and include Johnson (2011), Chatterjee and Ionescu (2013), and Stange (2012). Our analysis complements these papers in its focus on the role of skill premia on college enrollment and attainment.

1.1 Relation to Literature

Our analysis aims to provide implications for what one can expect, all else equal, from increases in the college premium (plausibly arising from the continuation of long-standing trends in skill-biased technological change) for (i) future aggregate human capital investment through college and (ii) consumption and earnings inequality. Importantly, we show selection on completion odds as well as other forms of risk matter for the answers.

Relative to existing work, our work is novel in being forward looking in the sense that we are interested in understanding the aggregate implications of further increases in the college premium and furthermore, in a setting where college is subject to noncompletion. While the question we pose is novel, the model we use is related to the work of Johnson (2011), Stange (2012), and Ionescu and Chatterjee (2013), while our

[1] The completion rate is the eight-year completion rate. Bound et al. (2009) stress that almost no completion occurs beyond this point. Additionally, while we will not address the issue, Heckman and LaFontaine (2010) describe an even greater barrier to college attainment: When properly measured, high school completion rates have been stagnant and lower than previously measured owing to the non equivalence of high school diplomas and GEDs.

[2] See, e.g., Restuccia and Urrutia (2004), Caucutt and Kumar (2003).

[3] Earlier seminal work of Altonji (1993) was perhaps one of the first analyses of human capital risk. Judd (2000) takes a portfolio-theoretic approach to the determination of the existence of excess returns to education.

[4] See also, Chen (2001, 2008) and Singh (2010).

model of the education prcoess itself is most closely related to ones developed in recent papers of Garriga and Keightley (2007), Abbott et al. (2012), Brown, Fang, and Gomes (2011), Hendricks and Leukhina (2012), Castex (2010a), and Johnson (2011). Garriga and Keightley (2007) and Abbott et al. (2013) both study higher education policy in settings that allow a rich array of general equilibrium feedback effects of large changes in higher education policy to operate.[5] However, aside from Garriga and Keightley (2007) these papers do not focus on the possibility of noncompletion and therefore abstract from college failure and earnings uncertainty. By contrast, as stressed above, our interest lies instead with measuring the strength of a given college premium on individual college enrollment decisions in the face of possible noncompletion and rate-of-return risk.[6] Castex (2010a) aims to account for enrollment over time in the United States but also abstracts from probabilistic noncompletion and rules out earnings risk as well. Johnson (2011) provides one experiment related to our main focus to predict enrollment under alternative skill premia but is otherwise not focused on the issue we study. Most recently, works of Restuccia and Vandenbroucke (2013) and Castro and Coen-Pirani (2013) aim to understand the sensitivity of educational attainment to its rewards, but as with the other work described above, they do not focus on the role of noncompletion and risk in determining the response of college investment to the college premium.[7]Rather, our focus is complementary to the latter– which provides a decomposition of the variety of forces influencing educational attainment in recent decades.

The two recent papers that do stress barriers to human capital acquisition and rate-of-return risk are those of Brown, Fang and Gomes (2011) and Hendricks and Leukhina (2011). The former calculates the present value of a college investment that has already been successfully completed and studies the role of taxation and social insurance policy in lowering the payoff to college, as well as insuring the risk of unemployment subsequent to college investment. However, Brown et al. (2011) do not focus, as we do, on completion, noncompletion, or the implications of the college premium for aggregate enrollment. With respect to college failure, Hendricks and Leukhina (2011) is relevant. Their paper has the goal of endogenizing dropout behavior and measuring the role of changing underlying ability amongst four-year enrollees over time in explaining the observed college premium (see also Hendricks, 2012; and Chatterjee and Ionescu, 2013). However, unlike our work, or that of Brown et al. (2011), Hendricks and Leukhina (2011) do not feature rate-of-return risk to human capital and do not feature assets markets to smoothing consumption.

The remainder of the paper is organized as follows. Section 2 lays out the model, and Section 3 describes the parameterization. In Section 4, we present results, first demonstrating the main result that college enrollment and attainment are not very responsive to increases in the college premium given current enrollment and risk levels. We then examine how noncompletion and risk affect college enrollment and attainment. College preparedness and financial resources are substitutes in the college enrollment decision, since a lack of preparedness lowers expected returns and increases risks, the latter can be self-insured with financial assets.

[5] See also Heckman, Lochner, and Taber (1998a,b) and Winter (2013).

[6] Specifically, endogenous skill prices would, for our purposes, introduce a needless additional step whereby we first would need to specify a production function that led to firms' demands for skilled and unskilled labor, and then we would need to vary TFP of skilled and unskilled labor in a way that delivered the various changes in the skill premium (and enrollment) that we were interested in at the outset.

[7] Restuccia and Vandenbroucke (2013) measure the extent to which skill-biased technological change can spur changes in educational attainment (relative to changes in perfectly forecasted, deterministic, life expectancy) They employ a model in which agents choose between three discrete education levels and solve for the equilibrium path of wages and attainment. Interestingly, they abstract from completion risk and changes in the ability distribution over time and find that their model generally overpredicts the level of attainment relative to what was observed, consistent with the idea of a previously nonbinding constraint becoming relevant. However, because there is no risk in educational investment, there is no notion of a change in the preparedness level of the marginal enrollee and no risk in the returns to human capital. Castro and Coen-Pirani (2013) proceed similarly but imposes a path for skill prices and study the evolution of educational attainment. Their paper also does not allow for college noncompletion and earnings risk.

This implies that college will attract either well-prepared and/or well-resourced students, while poorer and poorly prepared students are less likely to enroll. These results extend to the timing of the college premium, since delayed payoffs also dampen enrollment, especially for poorer students. Interestingly, higher high school wages can provide some insurance against drop out, but nonetheless, a temporary increase in the unskilled wage, such as a construction boom, dampens college enrollment and attainment, as has been found empirically. Finally, we document the increase in inequality that our model implies as the college premium rises, and Section 5 concludes.

2 Model Mechanism

Though our point is fundamentally a quantitative one, we turn first to a simplified model to convey the central forces at work.[8]

Consider a setting in which there is a large (continuum) of potential students who differ in their likelihood of failing to complete college , π. Assume that π is distributed uniformly on [0,1]. Let the cost of college be given by c and the earnings premium of college students (i.e., the earnings above that of non-college-educated persons) be given by p. Let $c < p$, ensuring that at least some agents will find college worthwhile. Assume next that all agents are risk-neutral, which allows us to define the (unique) "cutoff" or "marginal" student type as the one with noncompletion odds of π^* satisfying $c = (1 - \pi^*)p$. In this setting, the enrollment rate Ψ is simply the set types π for whom the expected benefit exceeds the cost: $(1 - \pi)p > c$:

$$\Psi = 1 - \frac{c}{p} = \pi^*$$

The possibility of noncompletion means that the "attainment rate" (the proportion of the population who finish college) is then:

$$A = \int_0^{\pi^*} (1 - \pi) d\pi$$

Which simplifies as follows:

$$= \pi^* - 0.5(\pi^*)^2$$

$$= \left(1 - \frac{c}{p}\right)\left(1 - 0.5\left(1 - \frac{c}{p}\right)\right)$$

$$= 0.5\left(1 - \left(\frac{c}{p}\right)^2\right)$$

The preceding implies that the proportion of enrollees who complete college, i.e., the "graduation" rate, G , is given by:

$$G = \frac{A}{\pi^*} = 0.5(1 + \frac{c}{p})$$

It is immediate that as the college premium rises, so do both attendance and graduation but at decreasing rates. That is, $\Psi(p)$ and $A(p)$ are increasing and concave. As for their limits, we see that as $p \to \infty$ (i.e., as the earnings premium for college rises without bound), $\Psi(p) \to 1$ (i.e., attendance becomes universal), which immediately implies that the graduation rate approaches the mean of the ability distribution: $A(p) \to 0.5$.

[8] We are extremely grateful to an anonymous referee for suggesting this approach.

What about college completion, G? We see from the expression above that it is decreasing in p, which reflects the declining mean ability of enrollees as the college premium rises.

The goal of this paper is to evaluate this mechanism quantitatively. In particular, we assess the likely position and shape of the mappings $\Psi(\cdot)$ and $A(\cdot)$ given empirically disciplined representations of preferences, college preparedness, college costs, earnings premia, and empirically plausible risk. As for risk, our quantitative model will assume a representation of preferences that reflects risk-aversion rather than the risk-neutral case above, and will build in the well-documented presence of significant uninsurable risk and inequality in initial wealth. Moreover, the model includes college noncompletion and earnings uncertainty. Taken together, these features allows us to generalize the argument above to quantitatively plausible results.

3 Quantitative Model

Our model features a large number (continuum) of individuals who differ, *ex-ante*, in their preparedness for college and in their initial wealth and who differ *ex-post* in their realized earnings, consumption, wealth, and educational attainment. Agents enter the model as high school graduates, and their first decision is whether or not to enroll in college, in either a two- or four-year program. Agents who enroll enter college immediately and progress toward a degree but understand that they may not complete it. Enrollees may also choose at any time to drop out. Once an individual exits college, or chooses not to enroll in the first place, they become workers who face earnings risk and have only a risk-free form of savings with which to smooth consumption. Workers thus face an entirely standard life-cycle consumption-savings problem. The aggregate quantity of human capital investment that is our focus will then be derived by aggregating enrollment decisions and collegiate outcomes over the distribution of initial preparedness and wealth.

3.1 Demographics and Preferences

There are three phases of decision-making. Households are modeled entering the economy as "*Young*" and will stay that way for K model periods. A period will be taken to be a year, and K measures the time between high school and successful college completion. College completion is defined to be the attainment of a bachelor's degree. Households then become "*Adults*" for J periods, which will be set to cover the length of time between college completion and retirement age, at which point they become "*Retirees.*" Young and Adult households order stochastic processes over consumption using a standard time-separable CRRA utility function with risk-aversion parameter α. All households discount the future exponentially at a common rate β. Expected discounted utility during pre-retirement life is thus: $E_0 \sum_{j=1}^{J} \beta^j \frac{c_j^{1-\alpha}}{1-\alpha}$.

Households value retirement resources according to a reduced-form "retirement felicity function" ϕ that is defined on wealth x_R taken into retirement, specified further below. This is a standard simplification (e.g., DeNardi, 2004) in models that build in a life cycle but are concerned mainly with early life decisions.[9] Like utility in working life, retirement felicity is modeled as a CRRA function of wealth and includes a weighting factor ν that will be calibrated to match wealth accumulation at the end of working life. Thus, discounted utility during retirement is given by $\phi(a_R) = \nu \frac{a_R^{1-\alpha}}{1-\alpha}$.

The general problem for the Young household is to choose consumption $\{c_k\}_{k=1}^{K}$ and make risky human capital investment (enrollment) decisions. Their enrollment decisions and the stochastic elements in their environment will leave them, eventually, with a human capital level $h \in \{HS, SC, C\}$, corresponding either to high school completion (HS), some college (SC), or college (C) attainment. To avoid clutter, we will

[9] It is particularly useful given our focus on the early-life decision problem of households who face a given skill premium and earnings and failure risk, as such decisions will remain insensitive to the temporally distant events of retirement.

suppress human capital in the notation below wherever it is obvious. Realized human capital conditional on enrolling will depend on the realization of uncertainty over college completion. Once Adults, households choose consumption $\{c_j\}_{j=1}^{J}$, and wealth during working life $\{a_j\}_{j=2}^{J}$, and wealth a_R with which they enter retirement.

3.2 Endowments and Market Arrangement

All agents are endowed in each period with one unit of time, which they supply inelastically in return for competitively determined wage. However, Young and Adult households face stochastic productivity shocks. These shocks are the source of uncertainty in the returnsex-post to even successful college completion.

Because households do not value leisure, they are modeled as simply receiving stochastic incomes in each period. The income process faced by households in the model is intended to represent precisely those risks that remain, *net* of (i) taxes, (ii) all private insurance mechanisms, and (iii) all *non*-means-tested public insurance programs such as the U.S. unemployment insurance system.

Log income is the sum of three components: an age (j) and human capital-specific (h) mean of log income μ_j^h, a persistent shock, z_j^h, and a transitory shock, u_j^h:[10]

$$\ln y_j^h = \mu_j^h + z_j^h + u_j^h \tag{1}$$

where

$$z_j^h = \rho^h z_{j-1} + \eta_j^h,\ \rho^h \leq 1,\ j \geq 2 \tag{2}$$

$$\ln u_j^h \sim i.i.d\ N(0, \sigma_{u,h}^2),\ \ln \eta_j^h \sim i.i.d.\ N(0, \sigma_{\eta,h}^2),\ u_j^h, \eta_j^h \text{ independent} \tag{3}$$

In addition, all households begin life as unskilled households, $h = HS$, and receive their initial realization of the persistent shock, z_1, from a distribution with an initial variance that differs from later ages. That is,

$$z_1 = \psi \tag{4}$$

where

$$\ln \psi \sim N(0, \sigma_\psi^2) \tag{5}$$

The income process can be interpreted as follows. To reflect heterogeneity prior to any direct exposure to labor market risk, households first draw a realization of the persistent shock z_1 from the random variable ψ with distribution $N(0, \sigma_\psi^2)$. In subsequent periods, households' nonasset income is determined as the sum of the the unconditional mean of log income μ_j^h, the innovation to the persistent shock η_j^h, and the transitory shock u_j^h. As is standard, the shocks to labor earnings during working age will depend on the human capital level of agents to reflect the fact that the risk characteristics of labor earnings appear to differ systematically by human capital level.[11] We note also that earnings shocks include, in an empirically disciplined manner, all the risks that even those who complete college still face, including the "underemployment" that people may experience following even a successful pathway in college.

As noted at the outset, ongoing discussions of enrollment and attainment often use current premia as a benchmark to gauge the responsiveness of enrollment. However, it is clear that current skill premia need not

[10] This is closely related to other standard specifications such as Huggett and Ventura (2000) and Storesletten et al. (2004), among others.

[11] See, e.g., Cagetti (2003), Chen (2001, 2008), Hubbard et al. (1994, 1995), and Storesletten, Telmer, and Yaron (2004).

have any bearing on college-related decisions, unless the current premium reflects the expected future college premium. Indeed, the currently observed college premium, precisely because it is historically high, perhaps should *not* to be expected to prevail over the working lives of current enrollees who complete college. We presume that households understand the future path of the *mean* college premium they can expect to receive at a given future age and that this average does not vary stochastically (though it can vary deterministically). Our approach is geared to shed light on the relative importance of skill premia that are *known, but possibly time-varying,* in the human capital investment decisions of heterogeneous households facing uncertainty.

Lastly, household risks are not fully insurable. Instead, households have access to a risk-free asset a that they may accumulate and which earns a constant (i.e., noncontingent) interest rate. Households may borrow as well, at a rate that includes a transactions-cost-related wedge on borrowing.

3.2.1 Means-Tested Transfer Income

Our model allows for means-tested transfers, $\varkappa(\cdot)$, represented as is standard (e.g., Hubbard et al., 1995) as a function of current age j, human capital h, net assets a_j, and income level $y_j^h = \exp(\mu_j^h + z_j^h + u_j^h)$. Empirically, these transfers may be sufficiently large for some households to alter decisions related to college investment. In the benchmark model, transfers will not depend explicitly on age and are specified as follows:

$$\varkappa(j, a_j, y_j) = \max\{0, \underline{\varkappa} - (\max(0, a_j) + y_j)\} \tag{6}$$

In the preceding, the term $\underline{\varkappa}$ refers to a societally-determined minimum floor on consumption that all agents are guaranteed. However, the transfers are means-tested, and an agent's assets a_j at age-j as well as their labor income y_j both count against these transfers one-for-one. But liabilities are not simply added to the agent's recources, thus requiring that the deduction itself be given by the sum of income y_j and the maximum of zero and the agent's assets a_j, i.e., $(\max(0, a_j) + y_j)$. Lastly, deduction of resources from the agent is meant to penalize asset holdings while not removing the guarantee of a floor on consumption and is thus restricted to be nonnegative. This gives rise to the outermost "max" operator, i.e., the least an agent can get by way of transfers is zero.

3.2.2 Retirement Income

A household's wealth level at retirement is then simply the sum of the household's personal savings a_{J+1} and the baseline retirement benefit $a_{\underline{\varkappa}}$. The amount $a_{\underline{\varkappa}}$ is the wealth level that, when annuitized at the discount rate R^f, and adjusted for the probability of survival for k periods, π_k, yields a flow of income each period equal to the societal minimum consumption floor $\underline{\varkappa}$ (identical to the consumption floor). The *lower bound* on retirement wealth $a_{\underline{\varkappa}}$ thus satisfies[12]:

$$\sum_{k=1}^{K} \frac{\pi_k \underline{\varkappa}}{(R^f)^k} = a_{\underline{\varkappa}} \tag{7}$$

3.3 Young Households and the College Investment Decision

As described above, there are K periods during which a decision maker is "Young." In the first period, individuals first draw income shocks y_j^{HS} from the distribution applicable to high school completers. This

[12]The reader may note that this specification does not allow for redistribution across agents in the manner that public retirement systems often provide. This abstraction is likely to be useful as it keeps the model simpler (i.e., do not have to track the entire history of earnings) but more importantly because we view it as reasonable to abstract from this feature of retirement plans given our focus on a much earlier-in-life decision (enroll or not in college).

informs them of the earnings they would receive if they decide not to enroll in college. Agents are given the choice to enroll in either two-year college or four-year college. All agents who enroll in two-year college have the option to transfer to four-year college after successful completion of the former. The central trade off between two- and four-year college is two-year college is cheaper, but leaves one less likely to complete four-year college should they succeed in the former and then transfer. For convenience, agents can only decide to begin enrollment at this stage, and not at later stages. As explained below, however, beginning college does not mean remaining there–exit is allowed and exit is, in some cases (as in the case of failure), required.

College of all types carries the possibility of non completion. Non-completion probabilities depend, in turn, on characteristics known to the agent: a measure of their ability and the type of college they are enrolled in. Households also draw a college non preparedness level, π, that gives the chance of failure in *the first year* of college.[13] If households do not enroll initially, they cannot enroll at a later date. In addition, college enrollees cannot work.[14] Households may borrow by using nondefaultable personal debt. Since the possibility of noncompletion and the risk aversion of individuals jointly make borrowing risky in terms of the utility of investment in college, access to credit alone may be insufficient to induce enrollment. We will show below that an enrollee's internal wealth position affects enrollment even after conditioning on failure risk and even when borrowing for college is allowed to the degree that college can be entirely funded by borrowing (no borrowing constraints).

Given knowledge of both the explicit costs of college, as well as the level of forgone earnings, households make the decision to enroll in college. If an individual enrolls, they must attend college while facing a risk, in each year, of failing to achieve satisfactory performance. Failure means that one is not allowed to continue in higher education. An enrollee's noncompletion risk evolves over time, where each year's failure risk is a fraction $\theta(\pi) \leq 1$ of the preceding year's risk, where $\theta(\pi)$ indicates that the probability of successful completion in each year depends on the initial failure risk. For example, if initial preparedness is given by π, failure risk from the current date's perspective is simply $\theta(\pi)\pi$ in year two, $\theta(\pi)^2\pi$ in year three, and so on to $\theta(\pi)^K\pi$. To reduce notational clutter, we suppress the dependence of θ on π wherever it is obvious.

If an enrollee learns that they have succeeded in a given year of college, they have the *option*–i.e., they are not required– to invest in additional years. Not all households who are informed of success will necessarily continue, however. Some will elect to leave college given a sufficiently high and persistent realization of income. We refer to such individuals as "dropouts." Those who voluntarily drop out and those failing to succeed both draw income from the shock process applicable to their human capital level. For those who have completed less than $k < \tau^{SC}$ years of college, the wage draw over life comes from the distribution of high school graduates agents and has mean $\{\mu_j^{HS}\}$. We define those with either a two-year degree or with $k \geq \tau^{SC}$ (acquired at a four-year college) as "some college," earnings will be drawn from the process with mean profile $\{\mu_j^{SC}\}$. The latter is not a proportional increase in earnings given the time relative to college completion, and it reflects the empirical regularity that college dropouts receive only a relatively small proportion of the income premium received, on average, by college graduates (e.g., Hungerford and Solon, 1989). After K periods (interpreted as years in the quantitative analysis) of successful completion, agents enter working life college educated and earn an expected payoff (in logs) denoted $\{\mu_j^C\}$.

For a household with currently low wealth and nontrivial failure chances, financing education with a fundamentally noncontingent instrument such as debt worsens outcomes under failure. Ex-ante, the distribution

[13] In practice, this individual-level assessment presumably arises from a variety of sources, including prominently a combination of family background, high school performance, and standardized test scores (e.g., Carneiro and Heckman, 2002; and Stinebrickner and Stinebrickner, 2008).

[14] This is relaxed in Section 6 on Robustness.

over future consumption (especially in the near term) induced by debt-financed college enrollment, *ceteris paribus*, makes college less attractive.[15] We will show that even without direct credit constraints, students do not always choose college even when the raw financial returns appear substantial. The additional risk induced with leverage has been emphasized recently (e.g., Cunningham and Santiago, 2008), and our model sheds some light on the magnitude of this effect.

As noted above, insurance markets against income risk are also incomplete, and all agents are instead endowed with only the ability to save in a risk-free asset that earns them return $1/q^f$. Agents may also borrow, but they must pay a proportional transactions cost on any debt they accumulate. To accommodate the presence of subsidized loans for the cost of college, we allow the wedge ζ to vary depending on how much the individual elects to borrow. Given access to subsidized college lending, borrowers face a net premium over the risk-free rate of ζ on any borrowing up to $\underline{a}_{Coll}$. For any borrowing beyond $\underline{a}_{Coll}$, but below the overall limit $\underline{a}$, an additional cost of intermediation (to reflect, for example, credit risk on personal lending) applies and leads to a net wedge of ζ_p. We therefore have:

$$q = \begin{cases} q^f \text{ if } a' > 0 \\ q^f - \zeta \text{ if } \underline{a}_{Coll} < a' \leq 0 \\ q^f - \zeta_p \text{ if } \underline{a} \leq a' < \underline{a}_{Coll} \end{cases}$$

As noted earlier, voluntary dropouts – leaving when one has not failed – are a possibility in our model. Since all agents receive a productivity draw in each period, even those who succeed in a given year of college may choose to drop out if their outside option is good enough.[16] While the model isolates genuine dropouts–those who leave but have the option to continue–from those who leave college as a result of either realized poor performance or anticipated poor performance, these are confounded in observed statistics on dropout rates. If an enrollee drops out prior to completing two years of college, he receives no premium, while if he drops out after τ periods, he receives a partial premium for completing “some college" (SC). To summarize, outcomes for human capital are given by:

$$h = \begin{cases} HS \text{ if no enrollment, or fail with } k < \tau \\ SC \text{ if enrollment eligible with } k \geq \tau \\ C \text{ if enrolled } \textit{and} \text{ no failure after } K \text{ periods} \end{cases}$$

4 Recursive Representation

The preceding is best described recursively. The state of any household can be expressed as follows. First, let π_k denote failure risk in the k-th year of college. Recall that a denotes household resources at the beginning of the period. Let k denote age while Young. For Young agents, the wealth level a should be thought of as the transfer that college-bound children expect to receive from their parents plus any internal funds they may have. Next, let z_k and u_k represent the persistent and transitory shocks to earnings, respectively. The state of a household is summarized by the vector $x = (k, a, z_k, u_k, h, \pi_k)$. To avoid clutter, in what follows we refer to the household state by x alone, with primes denoting one-period-ahead variables.

Three distinct value functions fully describe the household's problem. When Young, eligible households make the decision to enroll in postsecondary education by comparing the value of enrollment $V^E(x)$ with the value of not enrolling $V^{NE}(x)$. The value of enrollment in the first year (i.e., “*initial* college enrollment”)

[15] Chatterjee and Ionescu (2013) study the problem of how to insure against college failure risk, and in turn, show that an insurance program can increase enrollment rates substantially–suggesting that risk is indeed a relevant consideration in enrollment decisions.

[16] Current work of Lee and Shin (2012) also places emphasis on this.

is denoted by $V_1^Y(x)$. There are two kinds of college: community college and four-year college. In the first period, an agent who chooses to enroll must decide whether to enroll in community college or four-year college. Mechanically, this is done by comparing the value of community college $V_1^{CC}(x)$ with the value of enrolling in four-year college $V_1^{FY}(x)$. Substantively, the trade-off is this: an agent enrolling in community college will face lower tuition costs in the first two years than an enrollee at a four-year school but will also face a higher failure probability π, all else equal. We assume for convenience that if one does not enroll in the first period, one cannot enroll later on, and that the agent also cannot switch between four year school and community college in the second period. However, an agent who succesfully completes two years of community college may transfer to four-year college in the third year of college.

The choices above mean that the maximal utility attainable by a young agent eligible to enroll in college is given by the value function $V^S(\cdot)$ as follows. We use the superscript "S" in $V^S(\cdot)$ as a mnemonic to indicate a "successful" college student, i.e., one who has the *option* to enroll in college in the given period.

For Young agents in their first year, this requires that:

$$V^S(x) = \max(V^E(x), V^{NE}(x)) \tag{8}$$

where

$$V_1^E(x) = \max(V_1^{CC}(x), V_1^{FY}(x))$$

Trivially, in the first period of decision-making, all individuals, being successful high school completers by construction, have the option to enroll and are therefore classified as successful. If an individual enrolls, they understand that they will fail with probability π_k, in which case they will lose eligibility to continue in college and attain the conditional expected value available to nonenrollees, given their current persistent income shock z, $E_z V^{NE}(x')$. If they perform well enough to continue to the following year of instruction, something that occurs with probability $(1 - \pi_k)$, they realize an expected continuation value, given current persistent income risk z, of $E_z V^S(x')$.

The value of enrolling is then the solution to the following problem:

$$V^E(x) = \max \left[\frac{c^{1-\alpha} - 1}{1 - \alpha} + \beta \left(\pi_k E_z V^{NE}(x') + (1 - \pi_k) E_z V^S(x') \right) \right]$$

subject to the budget constraint if they enroll:

$$c + qa' + \Phi[1 - \gamma^{need}(x) - \gamma_k^{direct}] \leq a$$

$$a' > \underline{a}$$

and where

$$\pi_{k+1} = \theta(\pi)\pi_k,\ \theta < 1$$

In the budget constraint above, the term $\Phi > 0$ above denotes the annual cost of college, *prior to* all subsidies directly received by educational institutions from state, local, and federal sources. Direct subsidies in year-k of college are denoted γ_k^{direct} and apply to all enrollees. We will allow them to vary in order to reflect the fact that students have the option of initiating college education through community colleges, which are often substantially cheaper than even public four-year colleges. The term γ^{need} denotes further proportional reductions in the private cost of college arising from need-based aid.

Lastly, a denotes the wealth or resources available to an enrollee (in general, much of this will represent parental resources), with borrowing up to a limit $\underline{a} < 0$ allowed. This limit on total indebtedness is the sum of two specific borrowing limits. First, there is a limit on subsidized college loans given by $\underline{a}_{Coll} < 0$. Second, we allow for borrowing using other forms of personal credit, up to a limit of $\underline{a}_p < 0$. Thus, the total amount of credit individuals have access to is given by $\underline{a} = \underline{a}_{Coll} + \underline{a}_p$.

Because agents receive draws of earnings before deciding to enroll in college or not, and because these earnings would be foregone if they elected to stay in college, some may choose to drop out even when they have the option (by virtue of not failing) to continue. Under our maintained assumption that once one leaves college, one may not return, the continuation value of all those who exit college–whether by dropping out or failure–collapses to the following value function generated by the resulting standard consumption-savings problem:

$$V^{NE}(x) = \max \left[\frac{c^{1-\alpha}}{1-\alpha} + \beta E_{z,u} V^A(x') \right]$$

subject to the flow constraint:

$$c + qa' \leq a + I(\tau > k)y^{HS} + I(K > k \geq \tau)y^{SC}(x) + I(k \geq K)y^C$$

$$a' > \underline{a}$$

In the preceding, $V^A(\cdot)$ denotes the value of being an "Adult." Given the irreversibility of college nonenrollment, there is no difference between this value function and that applying to nonenrollees: Adults are, after all, nonenrollees. Thus, $V^A(\cdot) = V^{NE}(\cdot)$. Lastly, in the period immediately prior to retirement, households' optimal decisions satisfy:

$$V^A(x; j = J) = \max \left[\frac{c^{1-\alpha}}{1-\alpha} + \beta\phi(a_R) \right]$$

subject to the flow budget constraint

$$c + a_R \leq a + y^h(x) + \varkappa$$

$$a_R > 0$$

4.1 Aggregating Individual Decisions to Cohort Enrollment and Failure Rates

As clarified at the outset, our primary focus will be on understanding the *investment decision* of members of a cohort of young potential enrollees. To this end, we assign preference-, income-, ability- and education-cost-related parameters and then use equation (8) to solve for the household's optimal enrollment decision. Next, given the joint distribution of ability and income/resources, and remaining household state-variables, we can use equations (9) and (10) to immediately determine the behavior of aggregate college enrollment and attainment as a function of skill premia.

The flow of any new cohort of Young agents into college will depend on the joint distribution over the values of these state variables. Letting $\Gamma_1(x_1)$ denote the observed (cumulative) joint distribution of *age-1* Young households over the state vector, $x_{1,}$ (see Appendix A for details), $V^E(\cdot)$ and $V^{NE}(\cdot)$ the value functions associated with enrolling and not enrolling, respectively, and $I(\cdot)$, an indicator function over enrollment in college, we have that aggregate enrollment, denoted Ψ, is given by:

$$\Psi \equiv \int I(V^E(x_1) > V^{NE}(x_1))d\Gamma_1 \tag{9}$$

Similarly, given an underlying distribution of noncompletion probabilities as a function of the household's state, $f(\pi|x_1)$, the graduation rate is given by:

$$G \equiv 1 - \int f(\pi|x)I(V^E(x_1) > V^{NE}(x_1))d\Gamma_1 \tag{10}$$

Importantly, note that the noncompletion likelihood $f(\pi|x_1)$ is *endogenous*: It depends on the arrival of outside wage opportunities, household asset holdings, and place (in terms of years completed, for example) in the college regimen.

Our model of the enrollment decision is sufficiently rich to consider a realistic enrollment environment: It allows for the presence of initial uncertainty over collegiate preparedness, its gradual resolution over time, for exit and continuation decisions at a large number of dates (and hence the "option value" aspect to interim levels of education attainment), for stochastic and time-varying opportunity costs of college, for heterogeneity and nonindependence in individual resources and ability, for need-based aid and direct subsidies, for risky and uninsurable returns to any level of human capital that is successfully acquired, for a life-cycle consumption-savings dimension, and for credit market frictions in the form of a wedge on intermediation.

We next describe the parameterization of the model. We stress that our focus throughout is on individual decisions and the effects of prices, costs, and especially noncompletion and risk on them. In those cases where we allow for the cost of education to change, perhaps most naturally as the result of a change in higher education policies, our focus will remain on the impact of policy changes on individual decisions at the margin. That is, we provide measurements of the location and slope of the "supply curve" for college-educated labor. Our work therefore complements the empirical and policy-oriented literature (e.g., Garriga and Keightley, 2007; and Abbott et. al., 2013) by allowing for a large variety of simple counterfactuals, and emphasizing those in which college premia vary in isolation. One advantage of our approach is that we are able to use off-the-shelf parameter estimates nearly everywhere, with only a few specific parameters being calibrated to ensure that the model accounts for salient features of the data, specifically those related to enrollment, noncompletion, college premia, and wealth over the life cycle.

5 Parametrization

Aside from capturing the college investment decision, the problem we study is a textbook consumption-savings problem (e.g. Hubbard, Skinner, Zeldes, 1995). As a result, our model requires very few additional parameters: only those related to the probability of college completion and the valuation of wealth taken into retirement. All other parameters will be assigned values that are standard in the literature or values based on direct observation. The model requires us to assign values to four groups of parameters: those related to (i) preferences, (ii) education costs, (iii) familial resources (including credit availability) and collegiate preparedness, and (iv) stochastic processes for earnings as a function of educational attainment.[17] We present our parameter choices below in Table 1, followed by a detailed discussion leading to these choices.

[17] Specifically, for college, parameterization requires that we assign value to failure risk, $\{\pi^{(i)}\}_{i=1}^4$, noncompletion decay rates ($\theta(\pi)$), the retirement valuation parameter (ν), and earnings scaling factors (ϑ^C, ϑ^{HS}), and the wedge on intermediation costs (ζ).

Table 1: Parameter Values

Description	Parameter	Value	Source
(CRRA) Risk-Aversion	α	2.0	Standard
Discount rate	β	0.911	Calibrated
Valuation of wealth at retirement	ν	0.5	Calibrated
Skill-Premia: College/HS, Some-College/HS	$\{\vartheta^C, \vartheta^{SC}\}$	1.75, 1.18	Calibrated
Ann. drop in noncompletion prob. by math quartile	θ	{0.64, 0.48, 0.22, 0.02}	Calibrated
Initial noncompletion prob. by math quartile	$\{\pi^{(i)}\}_{i=1}^4$	{0.53, 0.51, 0.49, 0.16}	Calibrated
Borrowing cost wedges	ζ, ζ_p	0.03, 0.06	Calibrated
Borrowing limits for college, and personal credit	$\underline{a}_{Coll}$, $\underline{a}_p$	$-\$20,000$,$-\$20,000$	Calibrated
Direct college subsidy rate, share of college cost	$\gamma^{direct}_{benchmark}$	0.425	CK(2005)
Need-based aid: maximum, and income cutoff	γ^{need}	\$2,400, \$30,000	www.ed.gov
Familial Wealth Distribution	med_{a_0}, μ_{a_0}	\$3,000, \$11,000	GMV(2010)
Ability, Wealth Correlation	$corr$(test score, a_0)	0.3	C(2010b)

5.1 Additional Discussion of Model Parameterization

We now provide more detail on parameter choices and data sources.[18]

5.1.1 Preferences

First, with respect to preferences, there are only two parameters: the annual discount factor β, and risk-aversion α. Both β and α, though calibrated, take entirely standard values, at 0.945, and 2, respectively. For the valuation of resources taken into retirement, a_R, we use a simple CRRA specification with curvature α, $\nu \frac{a_R^{1-\alpha}}{1-\alpha}$, and impose the same value. We set $\nu = 2$ to ensure quantitatively appropriate wealth accumulation at retirement.[19]

5.1.2 Education Costs

College in our model represents all public higher education institutions, two- and four-year. This is the relevant set of institutions for three reasons. First, public entities account for the lion's share of enrollment (roughly 75 percent according to NCES 2000 data). Second, even though many will choose to attend more expensive schools, public higher education clearly remains a budget-feasible option for them. Third, public two-year colleges are cheaper than two years in public four-year colleges and allow for experimentation by enrollees by offering enrollees the option (e.g., Ozdagli and Trachter, 2012), conditional on successful completion, to continue to a four-year degree.

Out-of-Pocket Costs of College College costs in our model exclusively represent tuition and fees and do not include room and board. Room and board costs, as additional parts of college costs, are only relevant for those high school completers who would otherwise not have any costs of housing (e.g., they plan to continue living with parents). Our approach also helps ensure that we do not artificially limit enrollment by making the form of college in the model more expensive than the *cheapest* alternative available to qualified applicants seeking to attain a four-year college degree. Making college costs higher would make it easier to

[18] The reader interested in results may skip this section without loss of continuity.

[19] Naturally enough, ν turns out to be extremely *un*important for the question at hand, simply because retirement valuations are heavily discounted at the time of the college-enrollment decision.

establish widespread inframarginality at current enrollment rates. We therefore first specify college costs for the first τ years at the public two-year college rate, followed by the $K-\tau$ years at the public four-year college rate. The tremendous increase in two-year college enrollment relative to four-year college (Bound, Loevenheim, and Turner, 2010) suggests that we take this path because it carries the option to complete a four-year degree. As for the number of years it takes to earn the "some college" premium, we set τ=2. This is in line with the work of Ozdagli and Trachter (2011) and Hendricks and Leukhina (2011), for example, and can be interpreted as the payoff to a two-year degree.[20] Lastly, we set K at five years to represent the median time to college completion (NCES 2001).

In order to discipline the model's enrollment rates by test quartile, we calibrate the model to the same 1992 enrollment rates reported in Bound et al. (2009), and we use tuition and fee costs for that year from the College Board (2006) "Trends in College Pricing" Table 3a. This implies that households face tuition fees in current dollars, prior to any need-based aid, of approximately $7,000 per year to attend the last three years of public four-year college and roughly $2,500 per year for public two-year colleges. Because living expenses must be incurred irrespective of the college enrollment decision, we exclude room and board from direct costs of college.

Three parameters define education costs: the annual real resource cost of college Φ, γ_k^{direct} represents the average subsidy rate that is received by enrollees in the form of tuition and fee levels at public two- and four-year colleges, and $\gamma^{need}(x)$ determines need-based aid as a function of household type. We employ existing estimates for the direct subsidy to public four-year colleges in the range of 40 to 50 percent. Caucutt and Kumar (2003), for example, measure the subsidy at four-year public colleges and universities at 42.5 percent, which we will apply here for years $k>2$; this value is close to a more general consensus including Kane (2001), Table 14, for example, who suggests a number near 50 percent.

To parameterize need-based aid, $\gamma^{need}(x)$, we follow Clayton and Dynarski (2007) and the U.S. Department of Education and employ a simple linear function with two parameters governed by (i) maximal Pell grant of approximately $4,000 (in current dollars) is based on data in the year 1993-94 (Highlights of the Federal Pell Grant Program) and (ii) a constant reduction in Pell grants as a linear function of family resources, a_0, so that households with income greater than approximately $30,000 receive no aid.[21]

5.1.3 Familial Resources and Collegiate Preparedness

To parameterize the distribution of wealth available to potential enrollees, we employ a lognormal distribution of resources available to the student and therefore must assign values to only the mean and median of the distribution of initial wealth for Youths. The available wealth of enrollees will reflect not only their own private resources, if any, but also parental transfers. The latter, however, are not obviously proxied for by parental wealth since the willingness of parents to make such transfers is not directly observable. For the same reason, the level and covariance of familial resources available to *potential* enrollees (not just those who

[20] Because it is not central to our investigation, we have abstracted from heterogeneity among schools. We have parameterized the enrollment decision to a blend of two- and four-year public institutions. Public higher education enrolls the majority of college students (74 percent in recent NCES data), so we view this parameterization as capturing the cost structure facing the marginal student deciding whether or not to enroll in college. Those who enroll in more expensive schools face higher costs, so we assume they would surely have enrolled in the "cheaper" school that we model and hence this variation does not change our calculations of the enrollment rate. Nonetheless, the structure we employ could be used to model the distribution of students across schools, where the returns to attending various schools could vary along with their costs. Finally, as we noted, we study a problem in which households expect the skill premium to remain fixed. The important variation in this object seen in the past several decades means that large changes in the conditional return to college may occur from year to year. Expanding the model to allow for this added form of uncertainty is beyond the scope of this paper but seems worthwhile.

[21] Recently, the upper limit on Pell grants, in current dollars, has been increased above $5,000.

ultimately enroll) with any given test score is not well measured in the data. However, Kane (2001), Table 13, is informative; it finds that of those who report preparing financially for their children's college, only 25 percent with high school seniors had accumulated more than \$10,000. Relatedly, Gallipolli, et al. (2010) compute the distribution of inter-vivos transfers to the individuals between the ages of 16 and 22. We take the approximate midpoint of their estimates (see their Table 20) for the mean and median across parents of high school and college education, and we set the distribution of available resources to be lognormally distributed with a median, denoted med_{a_0} = \$3,000, and mean μ_{a_0}= \$11,000.[22]

Households may also borrow to finance college and consumption more generally. To set borrowing limits, we are guided first by the work of Carneiro and Heckman (2002, 2003), who argue that widespread borrowing constraints for education are implausible, and by the explicit set of guaranteed loan programs (the US government's Stafford and "PLUS" loan programs) to finance any amount in excess of the so-called "Expected Family Contribution."[23] We therefore set the debt limit to *always* allow a household to finance the entire cost of college (given the set of subsidies that are in place) and, to reflect current guaranteed loan programs, make the limits common across all households. Given the costs of college inclusive of all subsidies, we set this common borrowing limit at $\underline{a}_{Coll} = -\sum_k \Phi(1 - \gamma_k^{direct})$. In our benchmark economy, this amount is roughly \$20,000. In addition to this credit, we allow all households access to personal credit of roughly the same magnitude as well, i.e., $\underline{a}_p = -\$20,000$. Total borrowing capacity of households is therefore approximately $\underline{a} = -\$40,000$.

Because individuals are granted access to subsidized borrowing to cover the costs of college, none of our results arise directly from quantity constraints in credit markets. Still, credit use will interact with the uninsurable risks, especially noncompletion, as we emphasize. Bowen et al. (2009) suggest that "borrowing aversion" may play a role in the lack of a response in enrollment to the college premium even when credit availability is generous.

The interest rate associated with various asset positions are set as follows. First, the risk-free rate on savings is taken conservatively to be 2 percent (i.e., $1/q^f$=1.02). For borrowing, we allow for a wedge for intermediation to reflect any additional costs carried by borrowing relative to risk-free savings. However, to the extent that there is an additional subsidy to government lending for the financing of college, relative to that available for more general consumer credit in excess of the costs of college, we wish to allow for a separate wedge. In our baseline case, we impose a borrowing rate of 3 percentage points, i.e., $\zeta = 0.03$ up to the out-of-pocket cost of college. This is likely a lower bound on borrowing costs and is set to reflect access to subsidized college finance, as well as inter-family transfers (the current loan rate, for example, is

[22] These values are also similar to those from the NLSY documented by Johnson (2011), Figures 11 and 12. Also, in the estimates of Gallipolli et al. (2010), the median is remarkably stable, varying only from \$2,800 to \$3,500 when going from the least educated parents (all high school dropouts) to the most educated. The results turn out to be robust to substantial variations in the distributions of initial wealth, including ones with a mean as high as \$40,000 and a median of \$20,000. Higher values appear implausible. Gottschalck (2008), for example, using 2002 Census data reports median net worth, Table (4), for households between 35-44 at \$41,191 (\$9,512 excluding home equity) and that for 45-54 at \$82,435 (\$18,446 without home equity). But these measures, if used here, would be equivalent to presuming that *all* parental wealth is liquid and, furthermore, is available to college-bound households (or will eventually become available). Our measure is thus consistent with roughly one-half of these resources being essentially owned by the young enrollee.

[23] In recent work, Lochner and Monge-Naranjo (2011) argue that credit constraints may be more binding, but for the period that we choose as our benchmark, the work of Carneiro and Heckman (2001, 2002) seems decisive. See also Brown, Scholz, and Seshadri (2012). Relatedly, the work of Stinebrickner and Stinebrickner (2008) suggests that short-term credit constraints do not explain a substantial proportion of college failure. Overall, our findings are supportive of a point made forcefully by Carneiro and Heckman (2002): That the inability of children to "buy" the parental environs needed to make college a worthwhile investment is likely a key barrier, as opposed to the ability to borrow to finance an investment whose payoff to the marginal enrollee is well captured by the observed skill premium.

6.8 percent (nominal) for unsubsidized Stafford loans and somewhat higher for PLUS loans). For additional borrowing, we impose a wedge and set ζ_p=0.06, in line with estimates of Davis, Kubler, and Willen (2006), to reflect the costs of additional consumer debt incurred by the household. Thus, in the benchmark model, households can borrow to finance college at a real interest rate of 5 percent and can borrow for personal consumption at a real rate of 8 percent annually. The results turn out to be robust to variations in these wedges, as we show further below.

In our model, college completion conditional on enrollment is partially endogenous–due to the earnings opportunities students will receive while enrolled. However, the likelihood of successfully completing the first year of college is exogenous and given by the probability (1-π). We allow for four equally sized subgroups (quartiles) of the population in terms of their probability of completing the first year of college. This allows us to calibrate the probabilities $\{\pi_i\}_{i=1}^4$ to generate observed enrollment and noncompletion across quartiles of the test scores attained on the standardized mathematics test administered to all respondents in the NELS88 data.

Specifically, we follow Bound et al. (2009), who compute enrollment and noncompletion rates by performance on a standardized math test [24] and assign all individuals one of four failure levels, $\pi \in \{\pi^{(1)}, \pi^{(2)}, \pi^{(3)}, \pi^{(4)}\}$. In practice, the very low purely voluntary dropout rate in the model means we are severely restricted by the data in our choices for these four parameters, a clear virtue. This restriction can be seen easily. If (i) dropping out were either disallowed or if the arrival of outside opportunities was such that it was never optimal to drop out once enrolled and (ii) there was no decline over time in the likelihood of failure, then the probabilities we employed would be completely determined and exactly equal to the conditional probability of not completing college given one's position in the math test-score distribution.

Given the assigned values for "initial" noncompletion chances, we proceed by positing an underlying joint distribution of households over these values for preparedness (noncompletion risk) and their (log) financial resources and do not restrict them to be independent. To allow for dependence in a tractable manner, we assume that these two distributions are jointly bivariate normal and therefore must specify a single parameter to describe their dependence. In our benchmark model, we set $corr(\pi, a_0) = 0.3$, which is the benchmark value assumed in Gallipoli et al. (2010) and is similar to that measured by Castex (2010a, b) for the correlation between "Family Income" and "Ability" (as measured by standardized math test scores) in the NLSY79 and 97. The results are very robust to this parameter. As with several other parameters above, we view this approach as conservative. Assuming a higher correlation would once again create a larger population of inframarginal individuals: The rich would be even more disproportionately able and the poor disproportionately poorly prepared.

5.1.4 Earnings and College Premia

Our approach requires a position on the payoffs a potential enrollee should *expect* and the extent to which this is well-proxied for by the payoffs accruing to the current set of market participants across differing levels of education. But the observed college premium that prevails at any date, unless expected to persist over one's working life, is not the relevant fact for enrollment decisions. We target a path for earnings that generates a premium that approximates what prevailed over the period since 1993-2005, which, while in the spirit of rational expectations, does not presume perfect foresight. Our measures for earnings are derived from NLSY79 panel data using full-year earnings of those fully employed throughout the year.

[24] We stress that we do not directly calibrate overall enrollment rates; we calibrate the model to match the enrollment and completion, *conditional on standardized test scores*. The resulting aggregate enrollment rate closely matches the one estimated by Bound et al. (2009) and is slightly higher for the corresponding period than the one based on CPS data shown in Figure 1.

In our benchmark, we set the mean levels by age such that we match the targeted premium. That is, we locate two scalar coefficients, ϑ^C and ϑ^{SC}, on the mean of log earnings that yield targeted premia for college- and some-college-educated individuals: $\mu_j^{HS} = \vartheta^{HS}\mu_j^C$, and $\mu_j^{SC} = \vartheta^{SC}\mu_j^C$. In our benchmark setting, we set γ^C such that enrollees generate an average college premium of approximately 1.8 times that of high school completers, in line with Goldin and Katz (2006); that is, such that $\frac{E(y^C)}{E(y^{HS})} = 1.8$.

A second "price" an enrollee has to understand is the premium to completing "some college" and then failing to earn a degree, relative to the earnings they would receive as high school graduates. We set the premium expected by those who enroll, ϑ^{SC}, such that, after allowing for dropout decisions in the wake of good persistent outside options, we generate an observed premium of $\frac{E(y^{SC})}{E(y^{HS})}$ =1.15. This reflects an average over the various groups in our model who attain "some college" in line with the annual higher earnings premium for this group estimated by Kane and Rouse (1995), using data on those who attain two-year degrees. Importantly, in the benchmark economy, and unless otherwise specified, we hold the absolute earnings of unskilled households (i.e., high school completers) fixed, and we vary the payoff to college only. This reflects the overall flatness of real wages for unskilled households seen in recent decades and the rise in real compensation to those with college degrees. We will examine a case where this is relaxed.

To parameterize income processes across education groups, we follow standard estimates in the literature. Because our focus is on the role played by the return to human capital investments, tax policy matters. In particular, as Brown et al. (2011) stress, the progressivity of U.S. income taxes can lower the payoff to human capital because earnings to successful college completers tend to be higher and more compressed (temporally, because of the delay in generating earnings), while those of lower-skilled households are not only taxed at lower rates, but are also supplemented by other transfer programs. We use the estimates of Hubbard et al. (1994) as it provides estimates for earnings risk net of social insurance and does so for the same time period for which NELS data record noncompletion rates by standardized test scores.[25]

Following existing work (e.g., Caucutt and Kumar, 2003), we maintain the assumption that earnings, *conditional* on a given level of completed education, do not depend on initial ability (failure-risk). We stress that our parameter choices here, as elsewhere, are set to avoid simply forcing most into inframarginality with respect to college. Plausibly, failure risk (ability) and subsequent earnings conditional on educational attainment are likely to be negatively correlated, if at all. Lochner and Monje-Naranjo (2011), Table 1, for example, provide estimates that suggest that ability (as measured by standardized math test quartile) does influence average earnings over the life cycle, even conditional on college completion in the natural direction:

[25] Specifically, we use:

Parameter\Education Level	HS	Some College	College
σ_u^2	0.021	0.021	0.021
σ_η^2	0.025	0.025	0.014
σ_ψ^2	0.5	0.5	0.5
ρ	0.95	0.95	0.95

(11)

All the results here are robust to a much higher level of persistence, including $\rho = 0.99$. This is important because some (e.g., Storesletten, Telmer, and Yaron, 2004) advocate a near-unit root process (see Hryshko [2008] for a discussion of the evidence on persistence). We do not employ this in our benchmark model because it would only make the uninsurable risk faced by households still larger, and we wish to remain conservative along this dimension.

Lastly, the more recent rise in earnings volatility is partially muted by after-taxes and transfers, making the after-tax, after-transfer process of Hubbard et al. (1994) continue to offer a reasonable approximation to household level earnings uncertainty.

The lower one's failure risk, the higher the payoff conditional on completion.[26] [27] [28]

We remind the reader that our benchmark analysis aims at understanding the enrollment decisions of a given cohort in the face of a constant college premium that is expected to last with probability one over some or all of working life of the household. An alternative would be to go back in time to various dates, assume either that households understand the stochastic process governing skill premia over their lives or, more demandingly, have perfect foresight over the path that will unfold (see, e.g., Lee, 2005; Castro and Coen-Pirani, 2012), and compute enrollment rates. Given the documented longer-run variations seen in skill premia (e.g., Goldin and Katz, 2008) and the judgments that households must make regarding the future path premia associated with college documented in Kane (2001), our main focus is on how varying views on the path of premia should matter for enrollment.

6 Model Validation

We now examine the performance of the baseline model relative to the data. We start with the benchmark case–for which the model was calibrated. We then consider two time periods (early 1970s and mid-2000s) for which the model was *not* calibrated. The intent of this exercise is to verify whether the model of college enrollment and completion is robust and not simply fitted to the time period to which it is calibrated. In particular, we only change the parameters specifically governing the return to college, holding the rest of the structure of the model unchanged, and do not allow for other shocks to the macroeconomy. We show that the mechanism of the model is robust and sheds light on these other time periods, when the macroeconomy is subject to other shocks.

6.1 Fit of the Benchmark

Tables 2 and 3 document the fit of the benchmark parameterization along the two most salient dimensions. At the outset, note that our model studies high school graduates. Therefore, the enrollment rate (i.e., Ψ in our model) we report is that of the share of high school graduates who enter postsecondary education. Similarly, noncompletion rates (1-G in the model) are expressed as a fraction of enrollees, and our empirical measure of attainment (A in the model) is the percentage of young people (under 25) who have completed a four-year degree. In Table 2, we start with enrollment by preparedness level, where the "targets" are the measures reported by Bound et al., (2009), Figure 2, Panels C and A, respectively. The model is able to closely match enrollment and noncompletion rates across math test quartiles.

[26] See also Castex (2010b), Table E.2, and Hendricks and Leukhina (2011).

[27] In related work, Carneiro and Lee (2011) carefully measure the extent to which the quality of college students has fallen systematically since the 1960s and, strikingly, imply that the skill premium, if competitively determined in a period-by-period *spot* market for labor, would actually have been substantially higher were it not for the additional enrollees being of worse average quality. As a qualitative matter, such an effect would be expected to accompany changes in the premium for college completion, and indeed, do occur in our model.

[28] This is an issue studied by Hendricks and Schoellman (2012) as well. These authors find that this effect is important in influencing the observed college premium—making it substantially lower than if quality had not deteriorated with increased enrollment. These authors abstract from earnings risk by positing complete markets, and focus instead on the entirety of schooling decisions given a noisy signal of ability. Hendricks and Schoellman (2009) show that the price of skills has risen even faster than the change in wages, if one views the latter as a product of the price per unit of skilled labor and the level of skill possessed by a given worker. Lastly, see Hendricks and Leukhina (2011) for ongoing attempts to disentangle selection effects embedded in observed skill premia.

Table 2: Enrollment and Noncompletion by Math Test Quartiles: Model vs. Data

Quartile	Model (Ψ)	Data (Ψ)	Model (1-G)	Data (1-G)
1st	44%	45%	73%	80%
2nd	63%	62%	63%	63%
3rd	80%	76%	53%	48%
4th	92%	91%	30%	27%

Given that our model features both community (i.e. two-year) college as well as four-year colleges, we report in Table 3 the baseline model's performance on enrollment by institutional type, and note that it captures enrollment in each very well.

Table 3: Two- vs. Four- Year Enrollment: Model vs. Data

College Type	Model (Ψ)	Data
Two-year	22%	25%
Four-year	48%	48%
Total	70%	73%

Lastly, while omitted for brevity, the benchmark model replicates observed skill premia and also closely matches the paths of net worth over the life cycle for both high school educated and college-educated households as measured in the data. Taken as a whole, the benchmark economy generates decisions that match the data, at the level of aggregates related to college choices as well as of consumption and savings more generally.

6.2 Nontargeted Moments

We demonstrate above that the baseline model accurately captures enrollment and college completion across measures of (non) preparedness. These are the baseline results, however, we want to verify that the mechanism of the model is reliable more broadly. In particular, how well does the model do along dimensions it was *not* calibrated in any way to match? To answer this, we now evaluate the model's ability to account for college enrollment and completion across the past several decades. Specifically, we ask: allowing for changes in only net college costs and the college premium–are the model's implications for enrollment and noncompletion consistent with the data "out of sample" and able to demonstrate credibility of the fundamental structure and mechanism?

To explore this question, we evaluate the extent to which the baseline model can account for enrollment and noncompletion rates (again by mathematics test score quantile) for the high-school classes of 1972 and of 2004.[29]

[29] Interestingly, these authors also compute enrollment rates for the NELS72 cohort (who first enrolled in college in 1976) across the math test quartiles in those data and show that this group's enrollment rate was 48 percent (see their Table 1, Panel A) for that year, a time when the college premium was slightly above 1.5 (see again Figure 4). The model's predicted enrollment rate for this period is very close, at approximately 50 percent. Interestingly, the model-generated noncompletion rate (which is one minus the "Graduation Rate" in the figure) at the college premium of roughly 1.5 is also very close to the 49 percent noncompletion in the data for the NELS72 cohort (Bound et al. 2009, Table 1, Panel B). Of course, it is important to keep in mind that the model generates enrollment and completion rates based on premia that are expected to last through working life and, crucially, hold all other forces constant. It is not intended as a model of all dates in a particular segment of the higher education time-series landscape as, notably, is Castro and Coen-Pirani (2015). Nonetheless, this result does suggest that the model is broadly consistent with data, if households expected the college premium to be highly persistent.

This exercise requires additional data on enrollment, completion, college costs, and college earnings premia. Specifically, in addition to the NELS88 and NLSY79 data sets, we draw on two additional data sets. These are NELS waves that cover, respectively, educational choices and outcomes for the high-school graduating classes of 1972 (NELS72), and 2004 (NELS02). As with the benchmark, for income over the life cycle and measures of college readiness and collegiate outcomes over the life cycle in these two additional time periods, we employ the NLSY79 adjusted for the college premium prevailing in those periods. The exercise requires both the NELS and the NLSY since the latter does not allow us to distinguish between community- and four-year college enrollment (or completion).

We combine the two data sets to derive measures of enrollment (by college type), completion, and subsequent earnings over the life cycle as functions of these characteristics (as well as experience) and to then scale the estimated earnings processes so that they accurately represent earnings paths for cohorts before and after the baseline (HS class of 1992) cohort whose enrollment and completion by college type we use in our calibration of non completion chances. Moreover, for each of the three time periods for which we assess the model's performance, we adjust need-based and tuition costs to accurately reflect the distribution (across would-be enrollees from different households) of "net" college costs.[30]

Using these data sources, we now examine the predictions of the model over time. Table 4 below shows the model's implications for the high-school graduating cohort of 1972. The model closely captures both enrollment (denoted Ψ_{72}) and noncompletion (denoted 1-G_{72}) across preparedness quartiles (as measured by SAT-Math test quartile in the NELS72), as well as the proportions enrolling in two- and four-year colleges. The main gap between the model and the data is the overall enrollment rate, especially for those with the lowest test scores (and hence the lowest expected payoffs from college enrollment), which is about 15 percentage points higher in the data than predicted by the model. As emphasized in many discussions of college enrollement at the time (see e.g., Bruno [1984], and also Jones [2014, p. 211], for an informal argument), the explanation is surely the availability of deferment of military service associated with the draft imposed during the Vietnam War. Of course, this force is external to the model, yet interestingly, the mechanism of the model captures the rising pattern of enrollment across preparedness quartiles also observed in the data. And in line with the main prediction of the model, the increase in enrollment apparently due to the military service deferment is largest among those students least likely to enroll originally and smallest among those most likely to enroll. Hence, while the model is not calibrated to capture military deferment, the logic of the model is still evident in the response to this external force.

Table 4: Enrollment by Math Test Quartiles: Model vs. Data, HS Class of 1972

Quartile	Model (Ψ_{72})	Data (Ψ_{72})	Model (1-G_{72})	Data (1-G_{72})
1st	11%	22%	79%	74%
2nd	24%	38%	64%	64%
3rd	46%	56%	57%	52%
4th	73%	80%	37%	33%

Table 5 reports the model's predictions for the breakdown of two- and four-year college enrollees in 1972. The model predicts enrollment in two-year programs well, but the understatement of overall enrollment noted above manifests in disproportionately low four-year enrollment in the model compared to the data. This, again, seems plausibly connected with the pressures created by potential military service, which increased overall enrollment, especially in four-year schools, among students less likely to orginally enroll.

[30] The details are seen above in the subsection: "College and Earnings Premia."

Table 5: Two- vs. Four- Year Enrollment: Model vs. Data, HS Class of 1972

College Type	Model (Ψ_{72})	Data
Two-year	16%	15%
Four-year	23%	39%
Total	39%	54%

In sum, our baseline model appears to give sensible predictions for a period two decades prior to the one for which it was directly parameterized, given *only* the observed changes in college costs and college earnings premium. We now examine how the model performs for a period *after* the one for which the model was calibrated: the high school class of 2004. These results are presented in Table 6.

Table 6: Enrollment by Math Test Quartiles: Model vs. Data, HS Class of 2004

Quartile	Model (Ψ_{04})	Data (Ψ_{04})	Model ($1-G_{04}$)	Data ($1-G_{04}$)
1st	45%	51%	72%	80%
2nd	65%	67%	65%	66%
3rd	84%	80%	51%	44%
4th	94%	89%	27%	26%

The underprediction for enrollment seen for the Class of 1972 is no longer present, consistent with the hypothesis that high enrollment in the data was due to deferment of military service. The overall enrollment rate is very close to the data (see the last line of Table 7) and the rising pattern of enrollment rates with preparedness is evident in both the data and the model. Moreover, noncompletion rates (the last two columns of Table 6) fall systematically with preparedness, which is crucial for the selection mechanism emphasized in the model.

Lastly, Table 7 reports the implications of the model for the distribution of students between two- and four-year programs, compared to the data. In this case, we see that the model's overall enrollment matches the data well but has higher two-year enrollment and four-year enrollment that is somewhat lower than observed rates. This variation is interesting, given that the the only changes in model specification are those related to college costs and college earnings premia. We will examine below an experiment that would generate this change in the distibution of enrollment, driven by an increase in HS wages only. It has been argued in Charles et al. (2018) that the housing boom in the early 2000s generated just such an increase in the opportunties available to those with a high school education, which would reduce enrollment among two-year students, while at the same time providing additional financing capacity to homeowners, which would increase enrollment among four-year students (Amromin, Eberly, Mondragon, 2017). Hence, while the model matches overall enrollment well (72 compared to 74 percent), the mix of two- and four-year programs likely reflects changes in the housing market that we consider as a comparative static later in the paper.

Table 7: Two- vs. Four- Year Enrollment: Model vs. Data, HS Class of 2004

College Type	Model (Ψ_{04})	Data
Two-year	35%	27%
Four-year	37%	47%
Total	72%	74%

7 Results

As a whole, the performance of the model over a 30-year span suggests that it captures the central forces and the mechanisms that we hypothesize are driving college enrollment over time. The enrollment in two- and four-year programs match the data closely, and importantly, both enrollment and noncompletion rates track measures of preparedness as suggested by the mechanism in the model. We now employ the model to address the main question of this paper: What role does the U.S. college wage premium play in human capital investment when college is a risky, lumpy, and irreversible investment? Our main result is clear: The college premium does not drive enrollment and attainment beyond current levels given the risks faced by students. Next, by examining the determinants of the enrollment decision further, we can show how preparedness and financial resources substitute for one another. After all, well-prepared and well-resourced students face less risk and are better able to self-insure while poorer and less-well-prepared students are instead less likely to enroll. Moreover, the more delayed the payoff to education is, such as for graduate school, the less likely students are to enroll. Lastly, we will show that a temporary increase in the unskilled wage, such as a construction boom, deters college enrollment, even though it provides some insurance against the risk of noncompletion.

7.1 The Response of Enrollment and Attainment to Changes in the College Premium

We start with enrollment. Figure 1 shows how college enrollment responds to the college premium for degrees of college preparedness.[31] In general, enrollment rises with the college premium, but in the vicinity of the current college premium (above 1.5) the surface is relatively flat. That is, the response of enrollment to marginal changes in the college premium is small. This is exactly the sense in which our model confirms that most U.S. households are inframarginal with respect to the college premium.

As documented in, e.g., Goldin and Katz (2008), however, changes in the college premium over the past century have been substantial. Importantly, the calibration also shows that initial increases in the college premium above unity are associated with large increases in enrollment.

We turn next to the question of graduation rates which, by definition, are conditional on enrolling in college. The graduation rates mirror, almost exactly, the initial unconditional success probability of households. This is because for all but the lowest skill premia, students rarely voluntarily exit college. At any low college premium, however, the slight fall in graduation rates is due to enrollees lured away from college despite the successful completion of a given year of education. This is a useful implication of the model that, while beyond the scope of this paper, may benefit from further investigation: In the data, it will not always be easy to disentangle voluntary departures from involuntary-but-not-yet-realized failure. The model suggests that at current premia, essentially all exits will be failures rather than "voluntary" exits–even once one allows, as the model does, for the arrival of outside wage opportunities while enrolled in college.[32]

Given that voluntary "dropouts" are relatively rare, enrollment rates tend to drive college attainment for a given level of preparedness. As seen already, at a low college premium, overall enrollment is lower than currently observed. Figure 2 shows the stock of any entering cohort from the model that will eventually attain

[31] We stress again that our results throughout are for the case where enrollees expect the premium to last their entire working life (we will study a case further below with a temporary change in the skill premium). It is important to keep in mind, however, that the enrollment response is *asymmetric*.

[32] This is potentially relevant for understanding the findings of Stinebricker and Stinebricker (2008), who study the strength of credit constraints, which we do not allow for, and find that while they may push some to exit, the majority do not drop out due to binding credit constraints.

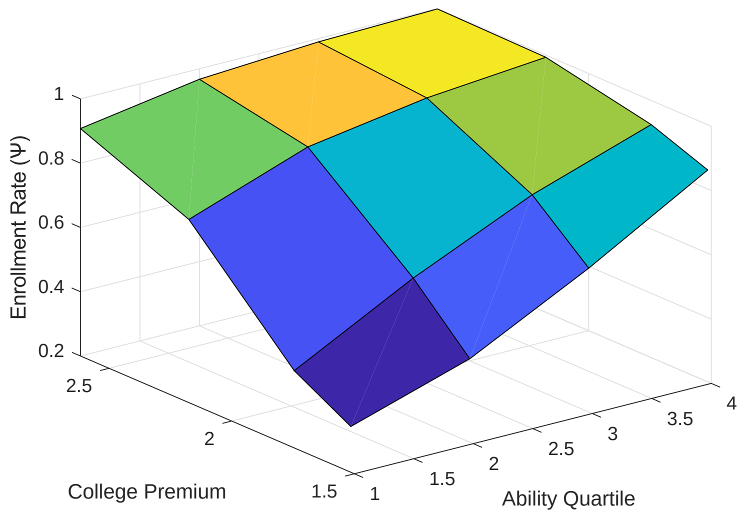

Figure 1: College Premium and Enrollment Rates Across Success Probabilites

a four-year college degree. The extremely low attainment rate among the very poorly prepared, especially at very low skill premia, reflects both the very low enrollment rate within this group documented in the previous figure and the low completion likelihood of the few who do enter. By contrast, the attainment rate is extremely high for very well-prepared individuals, who enroll *and* complete at very high rates. At very low college premia, attainment rates do rise with premia for the very well-prepared. This is a consequence of rising enrollment rates for this group.

7.1.1 The College Premium, Aggregate Enrollment, Graduation, and Attainment

Our results so far suggest that increases in the college premium, especially from levels close to that currently prevailing, should not, by themselves, necessarily be expected to significantly alter the stock of college-educated households. So far, though, we have displayed decision-making across college premia *conditional* on particular levels of preparedness. To derive specific implications from aggregate enrollment and attainment behavior shown thus far, we now use equations (9) and (10) to generate aggregate enrollment, graduation, and overall cohort attainment rates across skill premia. Figure 3 reports the results for the behavior of the aggregate enrollment rate, Ψ, the graduation rate, (1-Π), and college attainment. This demonstrates the main message of this paper: Increases in the college premium, all else equal, will not bring forth the college enrollees that they have in the past—the United States is now on the steep part of the supply curve of college-educated labor.

The model also suggests that under *a permanently* higher college premium (values of 2.1 or greater, e.g.) enrollment rates will rise only somewhat further than currently observed. Notice that graduation rates fall to the point where cohort attainment rates under these conditions are only slightly higher than their current

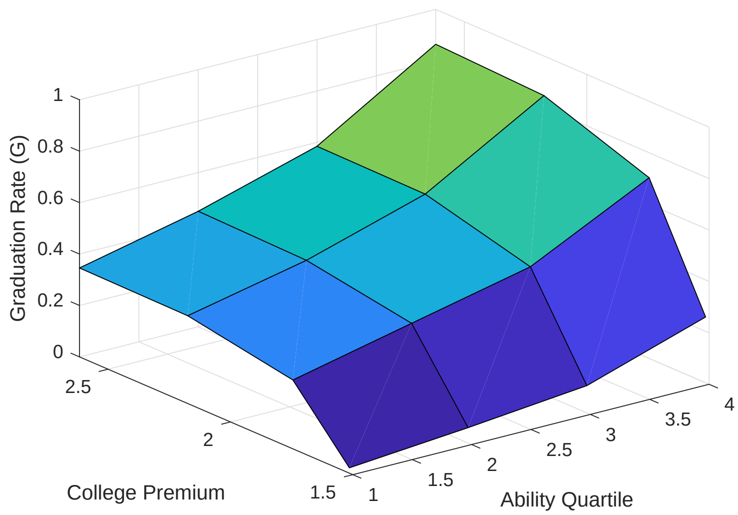

Figure 2: College Completion Rates Across College Premia and Success Probabilities

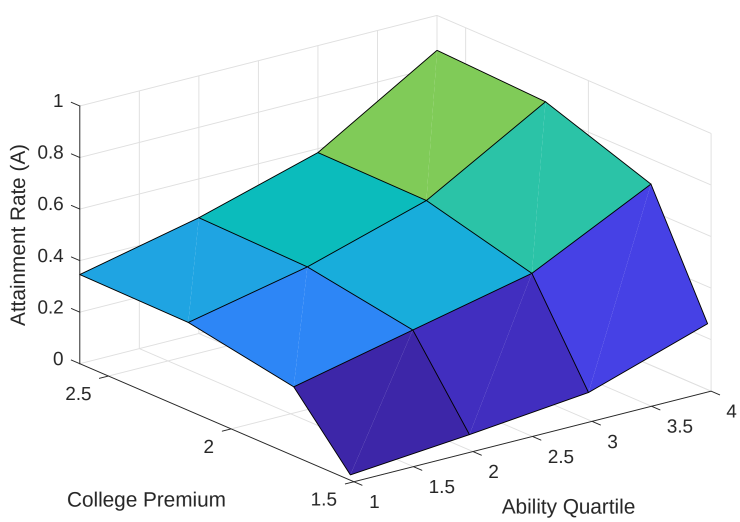

Figure 3: College Attainment: The Fraction of Each Cohort Completing College

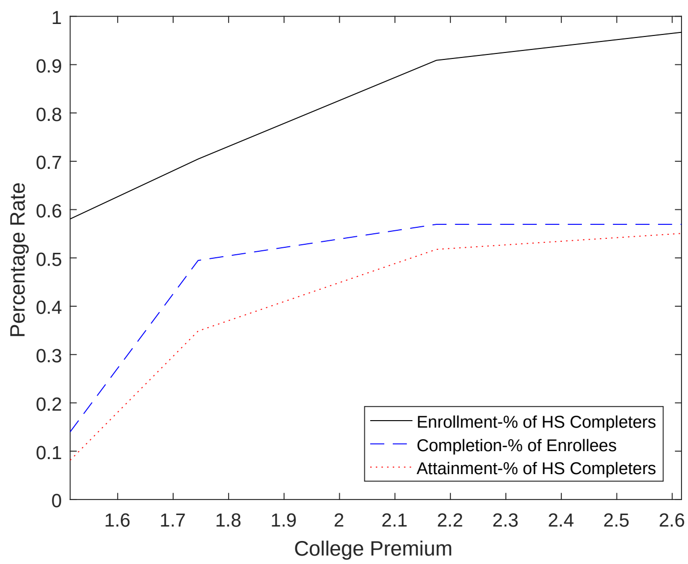

Figure 4: College Enrollment, Completion, and Overall Attainment Rates

level. This is because noncompletion rates rise as increasingly less well-prepared enrollees attempt college.

While enrollment is of clear interest, it is eventual completion, or the fraction within a cohort who have attained a complete college education, that represents skill acquisition. The attainment rate shown in the figure incorporates, by definition, both the initial likelihood of success and the enrollment rate of the group with that particular likelihood.

7.2 The Roles of Wealth and Preparedness in Enrollment Decisions

Our discussion above hints that enrollment decisions depend on factors beyond the college premium. We now examine directly the role of wealth, college preparedness, and opportunity costs in the enrollment decision. As a metric of the desirability of college, we calculate a *threshold level of wealth* necessary to lead a student to enroll in college. We solve for the level of assets, a_0^*, above which an individual chooses enrollment and below which she does not. We use this "wealth threshold" as a convenient metric for gauging the impact of changing conditions on the attractiveness of college.

Figure 5 shows that wealth thresholds vary systematically with skill premia for students with different characteristics. They are especially high for poorly prepared students, and especially so when skill premia are relatively low. This is intuitive: Wealth must be high for a student to attend college if the chance of noncompletion is high and the current college premium is low.

In contrast, for those likely to succeed, household financial resources have little impact on enrollment, irrespective of the college premium, while the opposite is true for those who are unlikely to complete college. Note, though, that because Figure 5 presents wealth thresholds that are conditional on completion likelihood, it says nothing by itself about how prevalent any particular preparation level is. However, it makes clear

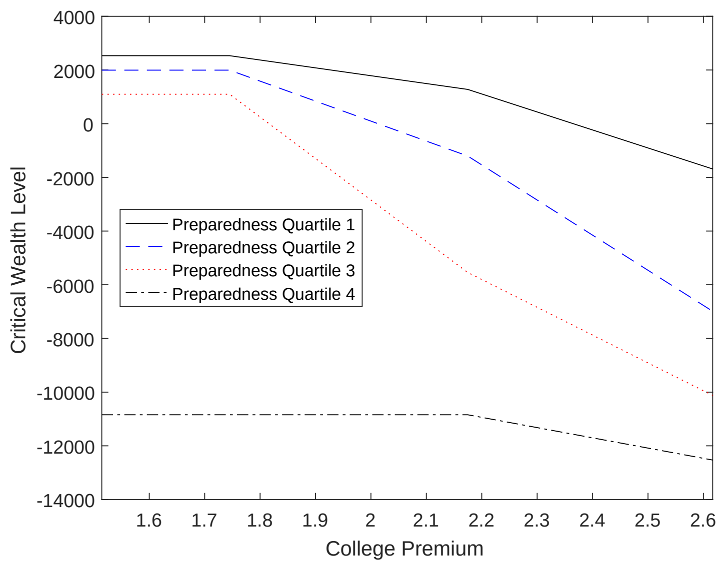

Figure 5: The Minimal Wealth Level Needed to Induce College Enrollment

that relatively poor enrollees will not, at current skill premia, be wealthy enough to find college attractive. Lastly, the very fact that there is a clear finite level of wealth above which the investment in college is worthwhile, and below which it is not, is clear evidence (at least within the model) that risk in human capital accumulation is playing a role.

The importance of risk is further clarified by observing the distribution of *consumption* that enrollees can expect to face as a function of their chosen educational attainment. Figure 6 shows selected moments of the present value of consumption up to each given age, under the benchmark college premium.[33] The dispersion in possible consumption outcomes is striking, as is the extent to which the distributions *overlap*; poor outcomes are possible along both trajectories, as are extraordinarily good ones. For example, while the highest trajectory is associated with college completion and the lowest with high school completion, the worst college trajectory is dominated by both the high school median and the high school 25th percentile. The possibility of such poor outcomes from "the college track" poses a substantial risk to poor and poorly prepared students, dampening college enrollment and attainment for these groups.

7.3 The Path of Future College Premia

Optimal investment behavior requires agents to forecast future prices and, in particular, future realizations of the college premium. Thus far, we have presented the model's predictions for changes in enrollment as a function of changes in the college premium that are expected to last an enrollee's entire working life. While

[33] The figure is calculated as follows. We first take all households in the last period of working life and locate the household whose present value places them in a given percentile. We then follow that household back through life in all earlier ages. In practice, at each date, we average over a small number (usually five histories) closest to the percentile under consideration and report the sample average of this group.

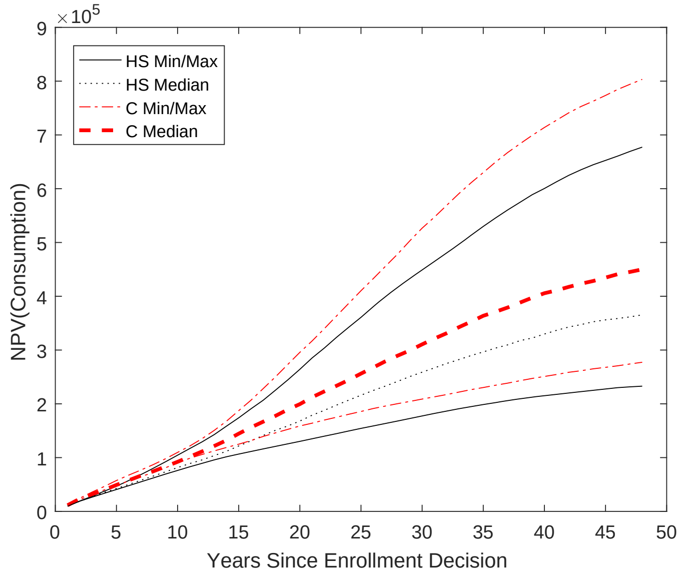

Figure 6: Consumption Outcomes for College Completers vs. Non-Enrollees

this is naatural given the question we pose, we note that it is also consistent with recent work by Donovan and Herrington (2014), and Castro and Coen-Pirani (2015), who both emphasize the relevance of "static expectations" in successfully accounting for changes in human capital accumulation in recent experience. But college premia have moved very substantially over the past 100 years (e.g., Goldin and Katz, 2008) and have moved substantially over periods much shorter than most individuals' working lives. Heckman and Navarro (2005) have argued that individuals may lack the ability to forecast a substantial portion of the payoffs to successful attainment of a given level of human capital.[34] And while our approach employs existing estimates of wage uncertainty, such measures do not allow for average premia to change as a function of age or experience. Even abstracting from uncertainty over future premia, therefore, it is of interest to understand the importance of the perceived *duration* of a given premium. The current premium is very close to historically high levels, and it is conceivable that it will drop in the future. In this section, we present results for how enrollment responds when the college premium moves up temporarily and then returns to its benchmark level (approximately 1.5, according to data of Goldin and Katz, 2008) at a date in the future that is expected by would-be enrollees.

We examine three cases in which we allow for premia to vary over time, presented below. In the first column, we examine the role played by premia that are expected to last for only the first decade of working life. In the second and third columns, we examine the strength of incentives coming from delayed premia of 10 and then 20 years from college completion.

Specifically, we evaluate the role of premia that depart temporarily from the historical mean of the college

[34] As to what people know about future payoffs from college, "Part V: Poor Information" in Kane (2001) is instructive. Among the findings of most interest were that students from more disadavantaged backgrounds seemed to systematically overestimate the costs of college.

premium over the period 1960-2010, which we measure at approximately 1.6. If that college premium were constant and permanent, our benchmark model gives an aggregate enrollment rate of roughly 61 percent. Each row specifies a departure in the college premium from 1.6, and each column specifies the years for which that departure lasts. For example, the top row of Table 8 examines the impact of an anticipated fall in college premia to levels below the historical average. The first column then shows how this affects enrollment when the deviation in premium occurs immediately after college graduation and lasts for 10 years. The effect on enrollment is substantial. Conversely, deviations to a higher premium (the bottom row) enhance enrollment if they arrive immediately but play little role in spurring enrollment if they occur only at later ages (e.g. bottom row, third column).

Table 8: College Premia, the Timing of Rewards, and Enrollment

$\frac{E(y^C)}{E(y^{HS})}$\Years of Altered College Premia	1-10	11-20	21-30
1.51	58.6%	60.3%	60.7%
1.75	67.1%	62.0%	61.2%
2.17	85.6%	65.7%	62.4%

The message is that if potential enrollees do not expect a given college premium to last for most of their working lives, or if they must wait for the premium to apply, enrollment will fall relative to current levels.

7.4 Movements in College Premia Arising from Changes in Wages for High School Completers

The analysis so far has varied the college premium by holding the payoff to high school completion fixed and changing the payoff to college completion. This, in general, is innocuous when it comes to understanding enrollment, as the *relative* payoffs to college versus non-college paths are of first order importance. However, in recent decades, a important proportion of the rise in the college premium has come from a systematic fall in the earnings of non-college-educated workers. This approach will not allow us to clearly understand college enrollment and attainment for shorter periods where the college premium changes as a result of an increase in demand for low-skilled labor inputs. For example, during the construction boom of the early and mid-2000s, Charles, Hurst, and Notowidigdo (2012) document an increase of unskilled wages along with an increased entry into manufacturing and construction-related employment, and they argue that it was a quantitatively important phenomenon in masking (initially) the secular fall in manufacturing employment. Because of the risk households face, a fall (rise) in the premium of a given level will have, in principle, different effects if it occurs from either a rise (fall) in the wages of the high school educated with college wages fixed or a rise in college wages with the high school wage fixed.

To examine the effect of the college premium on enrollment behavior, we consider two different experiments. In the first, we hold HS earnings constant and vary the college premium by adjusting college wages. In the second, we hold college earnings fixed and adjust HS earnings. We do the latter in order to examine a setting in which HS graduates see a change in their prospects – such as a construction or skilled craftsmen boom, as we discuss below – that does not change the absolute level of college earnings. The results are seen in Figure 7. Enrollment and attainment outcomes across skill premia where those premia are generated by holding fixed the college earnings path are denoted by "Fixed C earnings,"and those obtained by holding fixed high school earnings by "Fixed HS earnings." The latter is, of course, the baseline case. Throughout the exercise, we keep the absolute cost of college fixed so as to avoid adding confounding effects.

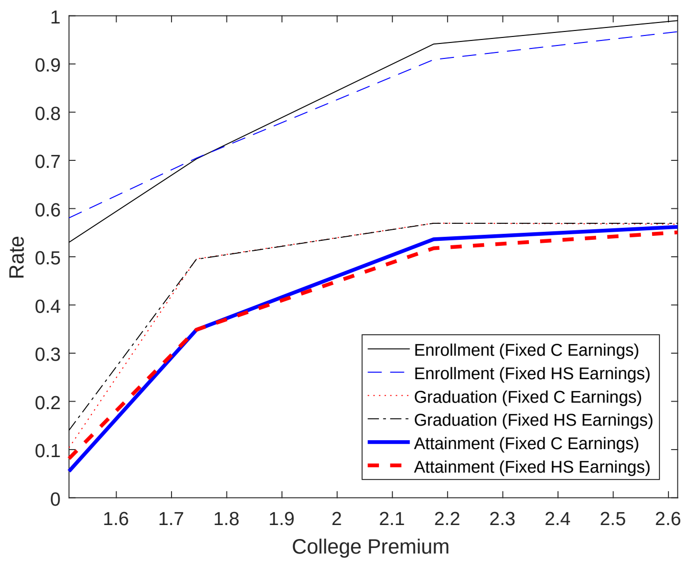

Figure 7: Attainment and Enrollment Across Premia: Benchmark Model vs. Fixed College Wages

Two points are immediate: First, enrollment, attainment, and graduation are largely similar across the two settings, demonstrating that the first-order effect of the college premium is the same in either case. This illustrates the importance, even in the face of risk, of the relative payoffs to college versus high school. Second, and where the two experiments diverge, is at college premia away from the benchmark. These differences in enrollment arise primarily from the risk of investment in a college education for a given college premium. At low premia, enrollment in college is actually higher when high school wages are high, because enrollment is less risky with more lucrative fallback employment. It is useful to note that this occurs even though the present value gap between college and high school is narrower at low premia than at high ones, something which, absent risk, should reduce college enrollment given the constant cost of enrollment. Alternatively, when high college premia come from low high school wages, college enrollment is less attractive, as failure could result in extremely poor consumption outcomes. Yet, it's again important to recognize that the fall in enrollment occurs despite the increased gap in the present value of earnings along each educational path, which, all else equal, should spur enrollment. These results emphasize that while the relative return to college (the premium) drives enrollment in college, the absolute level of resources (both wealth and HS wages) that potential students have to fall back on is important for enrollment decisions under uncertainty. These forces are also especially important for poor and poorly prepared students, as they face greater completion risk with less ability to self-insure.

As before, it is useful to measure the effect of a temporary compression of college premia arising from a rise in the wages of the unskilled. Table 9 presents the results. In this case, high school wages are modeled as "spiking" in a way that shrinks college premia for a temporary 10-year period of working life. In all other periods, earnings of the high school educated fall back to levels that, when averaged over the entire lifetime, yield a premium of 1.75. Our aim is to understand the likely implications of periods where labor market opportunities for unskilled workers improved disproportionately to their skilled counterparts. This type of path is aimed to characterize periods such as the construction boom of the late 1990s to the mid-2000s. The results are for the case where the college premium falls to a value of 1.51.

Table 9: College Premia, the Timing of Rewards, and Enrollment

Timing (years)	1-10	11-20	21-30
Enrollment Rate	36.1%	53.3%	63.7%

We see that enrollment shrinks substantially relative to the roughly 70 percent enrollment rate generated by the benchmark college premium of 1.75. Temporary spikes in premia at later ages also have strong deterrent effects on enrollment. The intuition here is that the ability to avoid both the explicit and opportunity costs of college while earning, as an unskilled agent, an amount similar to a college-educated one provide strong incentives to avoid college. Moreover, this occurs even though the increase in the absolute level of high school earnings makes the risk of investing in college more tolerable. In addition, a more subtle force is that a given level of compression in the college premium at later ages requires a larger increase in the absolute level of high school wages than at earlier ages. This is because college earnings paths feature greater returns to experience than unskilled earnings paths. Thus, despite their rarity, the model does suggest that episodes like the recent construction boom may well dissuade a large portion of would-be college enrollees from enrolling. Such a change will clearly have long-run implications for aggregate outcomes and may perhaps turn out to be one of the larger (ex-post) costs associated with the events of that period.[35]

7.5 Wage Risk, Noncompletion, and the Distribution of Consumption

Wage risk is present in our model for all skill levels. It is therefore tempting to conclude that it is in the nature of a "background risk" that does not alter the relative attractiveness of college versus high school educational attainment. This intuition misses the role of wealth and the high school wage in providing insurance, however. Wage risk creates the possibility of extremely poor outcomes ex-post, whereby households invest in college educations, fail to complete after spending substantial time, borrowing and forgoing earnings while there, and then receiving a persistent negative shock to wages. Absent wage risk, enrollees could at least count on being treated no worse than any other unskilled household with respect to wages. Given concavity, the absence of such a guarantee magnifies the risk of college investment, with the greatest risk borne by those with the lowest initial wealth. In Figure 8, we document percentiles of the consumption distribution across households who vary in their educational outcomes. The similarity of the distributions over consumption, despite the additional premium to earnings received by those with some college, makes clear that the costs of college can suppress consumption over the entire life cycle for those who fail to complete.

7.5.1 The Role of Market Incompleteness

To understand the role played by risk for enrollment in particular, as opposed to outcomes like wealth and consumption, it is useful to compare the baseline model with one in which markets are complete. That is, to what extent does uninsurable risk–even if it is an essential feature of the environment when it comes to the value of college–matter for the extensive margin decision of enrollment, and given enrollment, how it matters for college completion? Figure 9 below presents enrollment and completion for the complete-markets case.

[35] It should be noted that our model does not allow for workers to return to college later in life in order to take advantage of the college premium at later dates. As an empirical matter, this is not likely to be a severe restriction. See http://www.census.gov/hhes/school/data/cps/historical/TableA-6.pdf

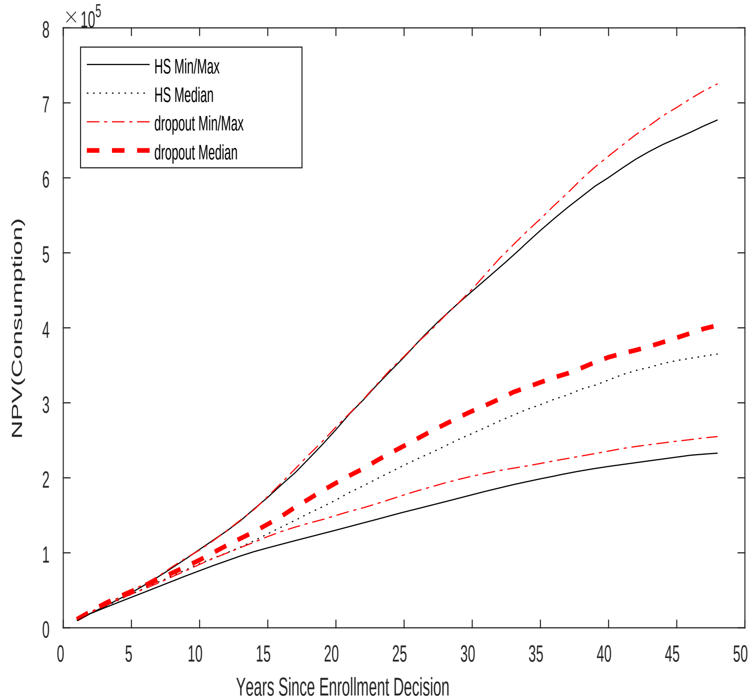

Figure 8: Consumption Outcomes for College Noncompleters vs. High School Graduates

Specifically, agents are assumed to be able to unload all wage risk (though not college completion risk) at actuarially fair rates and hence realize mean earnings (conditional on their educational attainment) with probability one.

We see that our results for college enrollment, graduation, and attainment are all virtually unchanged from the baseline model. If our goal were to make only the point that at the aggregate level the population of high-school completers are generally inframarginal to wage/earnings risk, we are arguably done. However, our goal is also to demonstrate that the college premium matters diferently across would-be enrollees from across the wealth and preparedness spectrum, and this requires a model in which allowance for market incompleteness is made.

7.6 The College Premium and Earnings Inequality

We noted at the outset that college premia and earnings inequality are likely to be related, and especially so if attainment is sluggish in the face of changes in premia, as it has been. To this point, we can calculate the response of earnings inequality to the college premium. Figure 10 reports the Gini coefficient of earnings for different values of the college premium and shows that inequality rises monotonically with the college premium. As we have stressed, college is a clear choice for the well-prepared, even at earnings premia that are modest relative to current ones. At high premia, enrollment is high, yet we have already seen that enrollment should not be expected to respond much to further increases in the college premium when starting from high levels. At the same time, the completion rate of additional enrollees is lower than the average. As a result, the human capital composition of the population does not change substantially, and the college premium feeds through directly to observed earnings inequality.

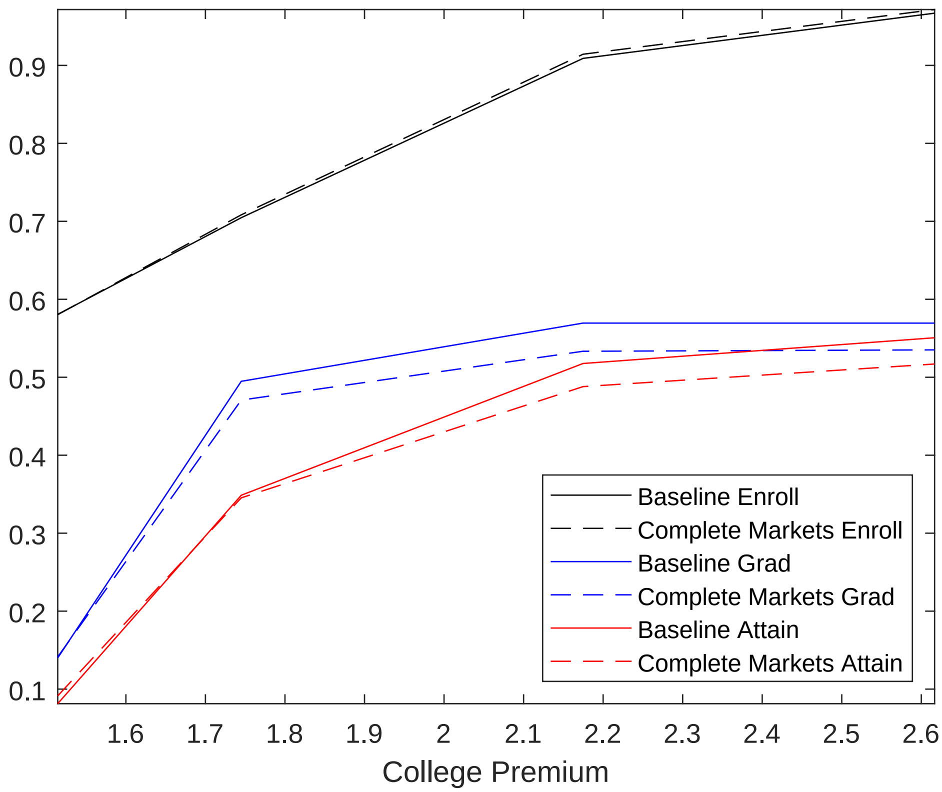

Figure 9: Attainment and Enrollment Across Premia: Benchmark Model vs. Complete Markets

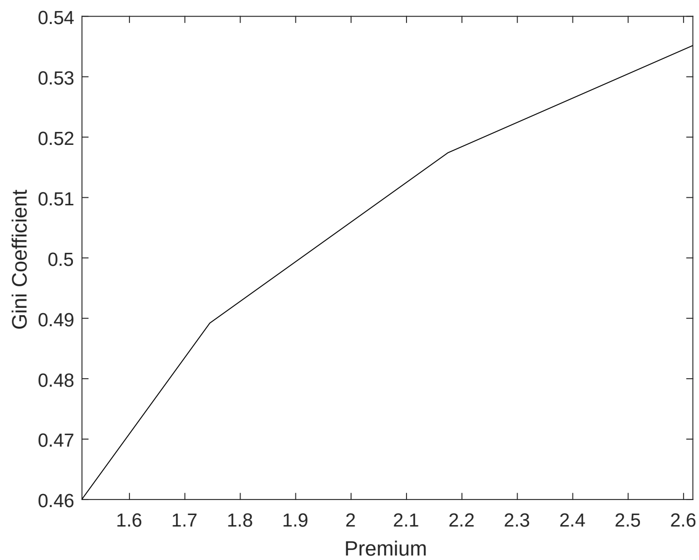

Figure 10: The College Premium and Earnings Inequality

The comovement of earnings inequality and the college premium has led U.S. policymakers to turn their focus to the higher education system as both a source of, and potential solution to, growing earnings disparities. Several recent proposals aim to expand access to college, typically by lowering its cost.[36] While it is beyond the scope of the current paper to provide a detailed analysis of such efforts, we note that our model suggests that, absent changes in the distribution of collegiate preparedness, college attainment will be difficult to increase. Indeed, the model suggests that broad-based reductions in college costs tend to flow to inframarginal households (similarly to increases in the college premium), which are those that are wealthiest and best prepared.

7.6.1 The Role of College Costs

We now briefly describe the implications of college costs on the sensitivity of enrollment and attainment to the college premium. This is potentially important as real college costs have risen in recent years (see, e.g., Jones and Yang [2016]), and it is relevant to know if the primary mechanism we focus on will, in the future, be substantially altered relative to our analysis. Note also that while we have already examined changes in college costs in the Model Validation section, those experiments also altered college premia to be consistent with data in different time periods. To focus entirely on college costs, all else equal, we increase the cost of college by 25 percent and leave all other parameters as in the baseline model. Moreover, in order to use this

[36] See, for example, the recent presidential proposals to increase access to college:
http://www.whitehouse.gov/a-better-bargain#education
and http://www.whitehouse.gov/blog/2015/01/08/president-proposes-make-community-college-free-responsible-students-2-years

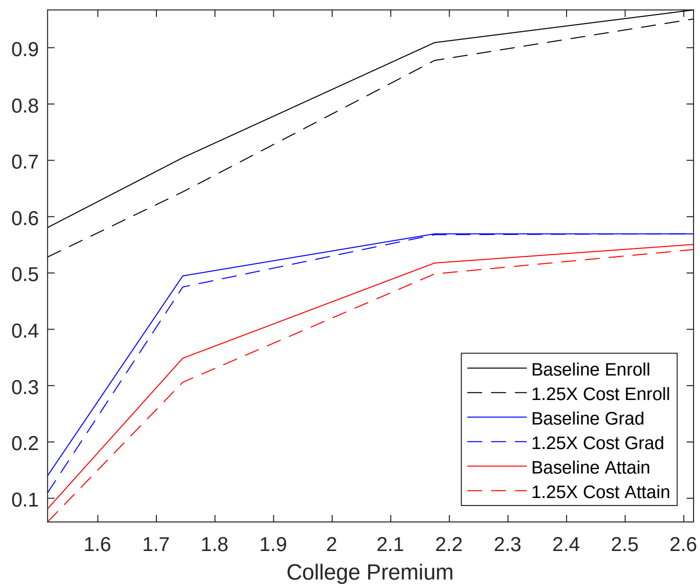

Figure 11: College Premia, Enrollment, and Attainment under Higher College Costs

exercise to tell about the relevance of college costs, we do not recalibrate the baseline model. In Appendix B, we conduct a related exercise, though in that case, we do recalibrate the model in order to advance the point that the model's main message is robust to alternative measurements of the cost of college. Figure 11 shows that our main findings are not changed by variation in college cost–other than by the natural implication that enrollment is lower at all college premia, with the effect being smallest at very high premia, where many potential enrollees are no longer deterred. To the extent that future college premia continue to move this direction, the effect of college costs on enrollment decisions (though not necessarily the "surplus" received by the would-be enrollee) are likely to be lower than they have been so far.

7.7 Robustness and Additional Discussion

Our results suggest that (i) the proportion of households who are currently marginal with respect to the college premium is small; and (ii) that the possibility of noncompletion and risks to subsequent earnings may both play important roles in college enrollment decisions. We refer the reader to Appendix B for a further detailed demonstration that the model's implications are robust, even though our model employs a variety of assumptions that bias it away from generating the widespread inframarginality we find.

Two caveats to our analysis apply. First, while rich along several dimensions, our decision model does not distinguish between various subgroups who, for whatever reason, face a college premium that is systematically lower *conditional* on college completion. For instance, if some minority groups face discrimination in the labor market and are generally less wealthy on average, they may well in practice respond strongly to expectations of a higher college premium.

Second, in keeping with our "ceteris paribus" approach, our model holds the preparedness of enrollees fixed and does not allow for changes in college preparedness that might arise in response to the college premium, such as via student effort in school and college, and parental investment early in the lives of children and adolescents. The work of Altonji, Bharadwaj, and Lange (2008a) goes directly to this point:

They focus on the question of how much the skills of the young have improved in response to the increase in the reward for skills. They find that while skills have improved across the entire distribution time, the bulk of those improvements have come from increases in the stock of parents with a college education. Our model suggests that the changes in college attainment seen from 1979 to 1997 are not likely to recur today. This makes the apparent inflexibility of skill acquisition among the young suggested by our model more likely to be a long-lasting phenomenon and one that may perpetuate the insensitivity of college enrollment to the college premium found in our investigation. Relatedly, James Heckman and coauthors have stressed in a variety of work (e.g., Heckman, Lochner, and Todd, 2008; and Heckman and Lafontaine, 2010) that the marginal returns can be very large already. Together, these findings are broadly consistent with our approach and findings for college enrollment and attainment at higher skill premia.

8 Concluding Remarks

In this paper, we have demonstrated that realistic risks present in college investment decisions will strongly limit the response of aggregate human capital investment to further increases in the U.S. college earnings premium. This means that under conditions reflecting those prevailing currently in the U.S., increases in the college premium can be expected primarily to increase earnings and income inequality rather than college attainment.

Our findings arise from a simple and intuitive mechanism rooted in the observation that neither college completion nor earnings are certain and, for many, *likely* to end in completion. The key idea is that the possibility of noncompletion generates asymmetric changes in the net return to college investment in response to a change in the college premium. Individuals in the model with relatively low failure risk see a relatively large increase in expected returns but are inframarginal under the current college premium. Current skill premia are more than ample to induce these individuals to enroll. Those with high failure risk, by contrast, see a much smaller increase in expected returns to college coming from a rise in the college premium. As a result, those currently choosing not to enroll remain largely inframarginal. The aggregate response of college investment to changes in the college premium is then governed by the behavior of those with more intermediate levels of completion likelihood. However, as a quantitative matter, our model suggests that this group too is inframarginal: Under the currently observed college premium, even those with relatively low completion rates enroll already. This implies that the remaining population of potential enrollees is now one that is (disproportionately) insensitive to the college premium. These potential enrollees are especially exposed to the risk factors identified in the model, arising from poor preparedness (completion risk) and earnings uncertainty (rate of return risk), magnified by low wealth (inability to insure) and leverage (borrowing)

Our findings have important implications: They strongly suggest that absent changes to college readiness, as especially strongly emphasized in the literature (e.g., Heckman and Carneiro, 2003; and the references therein), the continuation of long-running trends in the growth of demand for college-educated labor, (Goldin and Katz, 2008; and others) may lead to more rapid increases in income inequality than previously observed. As a corollary, our findings are consistent with slower future growth arising from a reduction in the ability to utilize skill-biased technological change. The preceding also being the larger body of work stressing the importance of early-childhood interventions (see, e.g., Heckman [2006] for a clear description of the rapidly declining payoff to late vs. early interventions, and the references therein). In this sense, our work, by taking some aspects of collegiate investment seriously–particularly noncompletion– simply underscores Heckman's point for the case of college investment.

The bigger point for our work that the important research into early-childhood education makes clear is that our findings must not be taken to mean that the sluggish response of aggregate skill predicted by our model is an inevitability. In fact, the combination of our finding that it may be sensible to take such a distribution as nearly invariant to the college premium conditions among those completing high school with that of the preceding literature finding large gains to early-life interventions, only further underscores the need to attend to the distribution of preparedness. The window for altering preparedness, and hence the potential for the US workforce to response to increase demand for skill, appears narrow.

References

[1] Abbott, Brant, G. Gallipoli, C. Meghir, and G. Violante. 2013. "Education Policy and Intergenerational Transfers in Equilibrium." Cowles Foundation Discussion Paper 1887.

[2] Akyol, A., and K. Athreya. 2005. "Risky Higher Education and Subsidies." *Journal of Economic Dynamics and Control* 29 (6): 979-1023.

[3] Altonji, J. 1993. "The Demand for and Return to Education When Education Outcomes are Uncertain." *Journal of Labor Economics* 11: 48-83.

[4] Altonji, Joseph G., Prashant Bharadwaj, and Fabian Lange. 2012. "Changes in the Characteristics of American Youth: Implications for Adult Outcomes." *Journal of Labor Economics* 30 (4): 738-828.

[5] Altonji, Joseph G., Prashant Bharadwaj, and Fabian Lange. 2008b. "The Anemic Response of Skill Investment to Skill Premium Growth." August 10. www.voxeu.org.

[6] Amromin, Gene, Janice Eberly, and John Mondragon. 2017. "The Housing Crisis and the Rise in Student Loans." Working paper.

[7] Arcidiacono, P, E. Aucejo, A. Maurel, and T. Ransom. 2016. "College Attrition and the Dynamics of Information Revelation." NBER Working Paper 22325.

[8] Autor, D., F. Levy, and R. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Explanation." *Quarterly Journal of Economics* 118 (4): 1279-1333.

[9] Bound, J., M. Loevenheim, and S. Turner. 2010. "Why Have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources." *American Economic Journal: Applied Economics* 2 (3): 129-157.

[10] Bowen, W., M. Chingos, and M. McPherson. 2009. *Crossing the Finish Line: Completing College in America's Public Universities*. Princeton University Press.

[11] Brown, M., John Karl Scholz, and Ananth Seshadri. 2012. "A New Test of Borrowing Constraints for Education." *Review of Economic Studies* 79 (2): 511-538. Oxford University Press.

[12] Brown, J., C. Fang, and F. Gomes. 2011. "Risk and Returns to Education: The Value of Human Capital." UIUC, LBS Working Paper.

[13] Bruno, Rosalind.. 1984, "Educational Attainment in the United States: March 1981 and 1980.", US Bureau of the Census, Current Population Reports, Series-P-20, No. 390.

[14] Cagetti, M. 2003. "Wealth Accumulation Over the Life Cycle and Precautionary Savings." *Journal of Business and Economic Statistics* 21 (3): 33, 9-353 .

[15] Cahuc, P., and A. Zylberberg. 2004. *Labor Economics*. MIT Press.

[16] Carneiro, P., and J. Heckman. 2002. "The Evidence on Credit Constraints in Post-Secondary Schooling." *The Economic Journal* 112 (482): 705-734.

[17] Carneiro, Pedro, and James J. Heckman. 2003. " Human Capital Policy." In *Inequality in America: What Role for Human Capital Policies?* edited by James J. Heckman, Alan B. Krueger and Benjamin M. Friedman. Cambridge, MA: MIT Press.

[18] Carneiro, P., and S. Lee. 2011. "Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000." *American Economic Review* 101 (6): 2309-49.

[19] Carneiro, P., K. Hansen, and J. Heckman. 2003. "Estimating Distributions of Treatment Effects With an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College Choice." *International Economic Review* 44 (2): 361-422.

[20] Carneiro, P., J. Heckman, and E.J Vytlacil. 2011. "Estimating Marginal Returns to Education." *American Economic Review* 101 (6): 2754-81.

[21] Castex, G. 2017. "College Risk and Return." *Review of Economic Dynamics* 26: 91-112.

[22] Castex, G. 2010b. "Accounting for Changes in College Attendance Profile: A Quantitative Life-Cycle Analysis." Bank of Chile Working Paper.

[23] Caucutt, E., and Kumar, K. 2003. "Higher Education Subsidies and Heterogeneity: A Dynamic Analysis." *Journal of Economic Dynamics and Control* 27 (8): 1459-1502.

[24] Charles, K., E. Hurst, and M. Notowidigdo. 2018. (forthcoming), "Manufacturing Busts, Housing Booms, and Labor Market Outcomes." *The Economic Journal*

[25] Chatterjee, S., and F. Ionescu. 2012. "Insuring Student Loans Against the Financial Risk of Failing to Complete College." *Quantitative Economics* 3 (3): 393-420.

[26] Chen, S. 2001. "Is Investing in College Education Risky?" mimeo, University of Rochester and SUNY Albany.

[27] Chen, S. 2008. "Estimating the Variance of Wages in the Presence of Selection and Unobserved Heterogeneity." *The Review of Economics and Statistics* 90 (2): 275-289.

[28] Clayton, J, and S. Dynarski. 2007. "College Grants on a Postcard: A Proposal for Simple and Predictable Student Aid." (with Susan M. Dynarski). Hamilton Project Discussion Paper 2007-01. The Brookings Institution, Washington, D.C.

[29] Cunningham, A., D. and Santiago. 2008. "Student Aversion to Borrowing: Who Borrows and Who Doesn't." *IHEP Report* March 4. www.ihep.org.

[30] Davis, S., F. Kubler, and P. Willen. 2006. "Borrowing Costs and the Demand for Equity over the Life Cycle." *Review of Economics and Statistics* 88 (2): 348-362.

[31] Dynarski, S. 2003. "Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion." *American Economic Review* 93 (1): 279-288.

[32] Fernald, J., and C. Jones, "The Future of U.S. Economic Growth," *American Economic Review Papers and Proceedings* 104(5): 44-49, May 2014

[33] Gallipoli, G., C. Meghir, and G. Violante. 2010. "Education Decisions and Policy Interventions: A General Equilibrium Evaluation." NYU Working Paper.

[34] Garriga, C., and M. Keightley. 2007. "A General Equilibrium Theory of College with Education Subsidies, In-School labor Supply, and Borrowing Constraints." mimeo, Federal Reserve Bank of St Louis.

[35] Goldin, C., and L. Katz. 2008. *The Race Between Education and Technology.* Cambridge, MA: Harvard University Press.

[36] Gottschalck, A.O. 2008. “Net Worth and the Assets of Households.” Census Bureau, 70-115. http://www.census.gov/prod/2008pubs/p70-115.pdf.

[37] Heckman, J. and R. LaFontaine. 2010. “The American High School Graduation Rate: Trends and Levels.” *Review of Economics and Statistics* 92 (2): 244-262.

[38] Heathcote, J., K. Storesletten, and G. Violante, 2009, “Quantitative Macroeconomics with Heterogeneous Households,” *Annual Review of Economics*, 1: 319-354

[39] Heckman, J. 2006, “Skill Formation and the Economics of Investing in Disadvantaged Children.” *Science*, 312.

[40] Heckman, J., L. Lochner, and C. Taber. 1998a. “Explaining Rising Wage Inequality: A Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents.” *Review of Economic Dynamics* 1 (1): 1-58.

[41] Heckman, J., L. Lochner, and C. Taber. 1998b. “General Equilibrium Treatment Effects: A Study of Tuition Policy.” *American Economic Review* 88 (2): 381-386.

[42] Heckman, J., L. Lochner, and P. Todd. 2008. “Earnings Functions and Rates of Return.” *Journal of Human Capital* 2 (1): 1-31.

[43] Heckman, J., and S. Navarro. 2005. “Separating Uncertainty from Heterogeneity in Life-Cycle Earnings.” *Oxford Economic Papers* 57 (2): 191-261

[44] Hendricks, Lutz. 2012. “The Evolution of U.S. Wages: Skill Prices versus Human Capital.” mimeo, http://www.lhendricks.org/Research/sojt/paper_sojt.pdf.

[45] Hendricks, Lutz, and Oksana Leukhina. 2014. “The Return to College: Selection Bias and Dropout Risk.” NBER Working Paper 4733.

[46] Hendricks, Lutz, and T. Shoellman. 2014. “Student Abilities During the Expansion of U.S. Education: 1950-2000.” *Journal of Monetary Economics* 63: 19-36.

[47] Hoxby, C. and C. Avery. 2013. “The Missing ‘One-Offs’: The Hidden Supply of High-Achieving, Low Income Students.” *Brookings Papers on Economic Activity* 1-50.

[48] Hubbard, G., J. Skinnner, and S. Zeldes. 1994. “The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving.” *Carnegie-Rochester Conference Series on Public Policy* 40 (1): 59-125.

[49] Hubbard, G., J. Skinnner, and S. Zeldes. 1995. “Precautionary Saving and Social Insurance.” *Journal of Political Economy* 103 (2): 360-99.

[50] Huggett, M., and G. Ventura. 2000. "Understanding Why High Income Households Save More Than Low Income Households." *Journal of Monetary Economics* 45 (2): 361-397.

[51] Hryshko, D. 2010. “RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles.” mimeo, University of Alberta.

[52] Hungerford, T., and G. Solon. 1987. “Sheepskin Effects in the Returns to Education.” *The Review of Economics and Statistics* 69 (1): 175-77.

[53] Johnson, M. 2013. “Borrowing Constraints, College Enrollment, and Delayed Entry.” *Journal of Labor Economics* 31 (4): 669-725.

[54] Judd, K. 2000. “Is Education as Good as Gold? A Portfolio Analysis of Human capital Investment.” mimeo, Hoover Institution.

[55] Jones, Charles, I., Macroeconomics, 3rd Edition. WW Norton.

[56] Kane, T. 2001. “College-Going and Inequality: A Literature Review.” Russell Sage Foundation Working Papers.

[57] Kane, Thomas and Cecilia Rouse. 1995. “Labor Market Returns to Two - and FourYear - College.” *The American Economic Review* 85 (3): 600-614.

[58] Lee, S. Y., and Yongseok Shin. 2013. "The Option Value of Human Capital," mimeo, http://lee.vwl.uni-mannheim.de/materials/hc_option.pdf.

[59] Layard, R., and G. Psacharapoulos. 1974. “The Screening Hypothesis and the Returns to Education.” *Journal of Political Economy* 82 (5): 985-98.

[60] Lochner, L., and A. Monge-Naranjo. 2011. “The Nature of Credit Constraints and Human Capital.” *American Economic Review* 101 (6): 2487-2529.

[61] Restuccia, D., and C. Urrutia. 2004. "Intergenerational Persistence of Earnings: The Role of Early and College Education." *American Economic Review* 94 (5): 1354-1378.

[62] Restuccia, D., and G. Vandenbroucke. “The Evolution of Education: A Macroeconomic Analysis.” *International Economic Review* 54 (3): 915-936.

[63] Ozdagli, A., and N. Trachter. 2011. “On the Distribution of College Dropouts: Household Wealth and Uninsurable Idiosyncratic Risk.” Federal Reserve Bank of Boston Working Paper 11-8.

[64] Singh, A. 2010. “Human Capital Risk in Life-Cycle Economies.” *Journal of Monetary Economics* 57 (6): 729-738.

[65] Stange, Kevin. 2012. “An Empirical Examination of the Option Value of College Enrollment.” *American Economic Journal: Applied Economics* 4 (1): 49-84.

[66] Stinebrickner, Todd, and R. Stinebrickner. 2008. “The Effect of Credit Constraints on the College Drop-Out Decision: A Direct Approach Using a New Panel Study." *American Economic Review* 98 (5): 2163-84.

[67] Storesletten, K., C. Telmer, and A. Yaron. 2004. "Consumption and Risk-Sharing over the Life Cycle." *Journal of Monetary Economics* 51 (3): 609-633.

[68] Winter, Christophe. 2014. "Accounting for the Changing Role of Family Income in Determining College Entry." *Scandinavian Journal of Economics* 116(4): 909-963.

9 Appendix: Robustness

In this Appendix, we show in detail that our main conclusions are strongly robust with respect to a wide array of *quantitatively relevant* alternative environments. We investigate, in turn, assumptions on credit markets, on college costs, and on familial resources. Importantly, we remind the reader that in each case where we deviate from the benchmark parameterization, we *recalibrate* model parameters to locate the closest match we can to the targets of the model in terms of enrollment and noncompletion by mathematics test quartile. The results are given in Figures 12 and 13. Once we have recalibrated the model, we proceed as usual and evaluate the sensitivity of enrollment and completion to the college premium. Recall that ensuring robustness of our main results requires that we *do not overstate* the fractions who are either extremely well-prepared, or extremely poorly prepared, since such households will be inframarginal in general.

For robustness to the assumed structure credit market, we tightened credit constraints by 25 percent (denoted "Tighter Credit Constraint") and then eliminated the wedge on borrowing costs ζ (i.e., we set $\zeta = 0$ and denote this as the "Frictionless Borrowing" case). For robustness with respect to the cost of college, we relax the benchmark model's assumption that all households have access to the path whereby they enroll and pay the two-year college rate for two years, after which they faced the four-year rate (public, in-state) for the rest of college. Instead, we assume that the pathway from community college to four-year college is no longer available. Thus, all students must pay the four-year rate in all years of college. This is labeled "No Community College."In addition to this, we examine a variant of the model in which we increase by 25 percent the cost of college ("1.25X Cost"). Next, to validate robustness with respect to inequality in familial resources, we double the skewness of initial wealth ("2 x Mean/Median a_0") and also present robustness to the correlation between wealth and math scores by doubling it from the benchmark case to 0.6 ("2 x Test Score/a_0 Correlation"). Full details on the parameters used in each of these exercises are available upon request.

In each of the cases described above, we focus on the two central figures related to college enrollment and attainment, and we see that in every one of the alternatives, outcomes are essentially the same as in the benchmark model: Enrollment and attainment are responsive to premia lower than ones currently observed and then become sharply less responsive at premia near, or higher than, current levels.

Next, we display the behavior of attainment across different values of the college premium and see that it also remains very robust across the large variety of alternatives we consider.

We comment next on the use of a discrete number of types of agents, particularly with respect to failure risk. For the question of how much enrollment and attainment will respond to changes in the college premium, the issue, in the language of a standard Roy model, is the distribution of "surplus" or net gain from enrolling in college across the population. That is, at any given college premium, the response of enrollment and attainment in the aggregate depends on the proportion of potential enrollees located in a region where they are close to changing their enrollment decision. How much inframarginality is predetermined by restricting attention to four discrete failure-probability groups? To address this, we have conducted robustness exercises in which we have allowed for more categories but required that when aggregated into quartile-based enrollment and noncompletion, we match the original targets. The results are essentially identical. In one sense, this is not surprising because there are already two features of the model that make its implications not very sensitive to the inclusion of more failure groups. Specifically, agents differ in two other ways that are more "continuously" distributed in the model, wage opportunities and wealth, and this guarantees more richness and heterogeneity.[37] Intuitively, at very low college premia, enrollment is overwhelmingly from the

[37] To be clear, in our computations, we are always restricted to discrete sets of values for all variables, including wealth and wages. But these variables have relatively very fine discretizations and hence introduce ample heterogeneity.

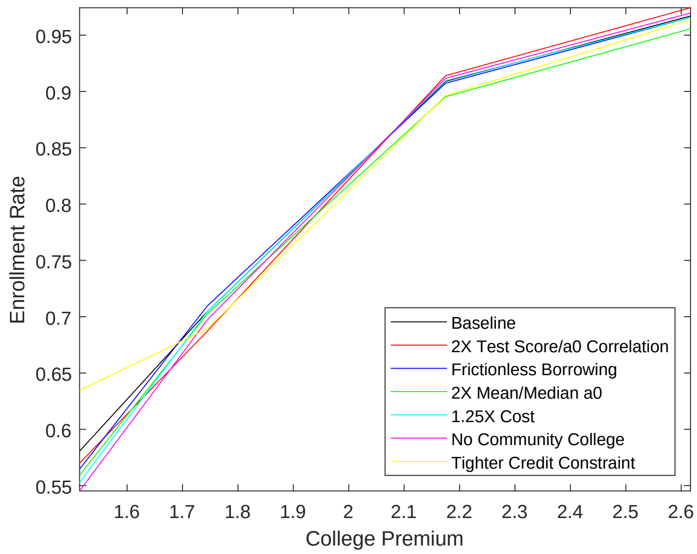

Figure 12: Robustness: Skill Premia and Enrollment Rates

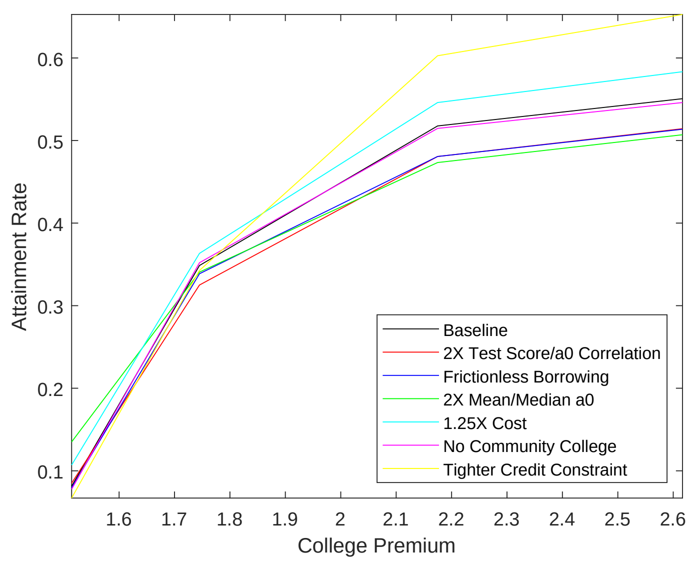

Figure 13: Robustness: Skill Premia and Attainment Rates

wealth-richest, wage-poorest, *and* best-prepared subset of the population. And a large fraction of these households complete college (i.e., do not fail and do not drop out voluntarily). As the college premium rises, less wealthy households with the group of highly prepared students begin to enter. As premia climb further, the wealth-poorest, wage-richest among the best prepared and the wealth-richest, wage-poorest among those in *next* quartile of completion chances both enroll, and so on. This process makes enrollment and attainment move *smoothly* upward, though with a systematically declining rate as the college premium rises–which is our main point.

9.1 An Extension of the Baseline Model to Allow for Effort in College

A final, more complex, robustness exercise is allowing for the possibility that noncompletion odds can be altered through effort. To remind the reader, in the baseline model we have used so far, as well as in all the robustness variations we have so far considered, an agent cannot endogenously influence their probability of success upon enrolling in school. Instead, they take it as given (i.e., as an immutable function of their exogenously determined preparednes, as measured in the data by the quartile of mathematics standardized test score to which they belong). That restriction allowed us to cleanly focus on the the effects of skill premia changes on enrollment across ability levels.

However, for completeness and robustness, we now examine the case that allows for endogenous choice of effort level. In particular, it is important to gauge the extent to which in a high-college-premium setting, enrollees' compeletion incentives would lead them to exert greater effort to ensure college completion. In our analysis, this simply means that it is of interest to allow for greater effort when the college premium is high than when it is low.

For simplicity, we study a model in which we retain all the features of the baseline setting, but in addition allow for a binary effort choice. Specifically, in each period, an agent who enrolls in school can chose to exert low effort L or high effort H. If they exert high effort, they incur disutility of effort d_e but will have a lower failure probability in that period. Specifically, college enrollees can now reduce their failure probability (which is as before dependent on ability quartile) by a scalar $s < 1$ by exerting high effort relative to what we would obtain under low effort. We calibrate d_e and s to our benchmark case and then allow the college premium to vary. Formally, the college enrollment decision problem now is modified to reflect the role of effort as follows:

$$V^E = \max\{V^E_L(x), V^E_H(x)\}$$

$$V^E(x) = \max_{c,e}[\frac{(c^\lambda(1-d_e)^{1-\lambda})^{1-\alpha}}{1-\alpha} + \beta(\pi_e EV^{NE}(x') + (1-\pi_e)EV^S(x')]$$

$$e \in \{L, H\}$$

$$1 > d_H > d_L = 0$$

$$V^{NE}(x) = \max_c[\frac{c^{1-\alpha}}{1-\alpha} + \beta EV^A(x')]$$

The key change is seen in two places: (i) in within-period preferences, which now reflects the cost imposed by effort, via the multiplicative term $(1-d_e)^{1-\lambda}$, and (ii) the implicit dependence of noncompletion odds, π_e, on effort. All else remains as in the baseline economy.

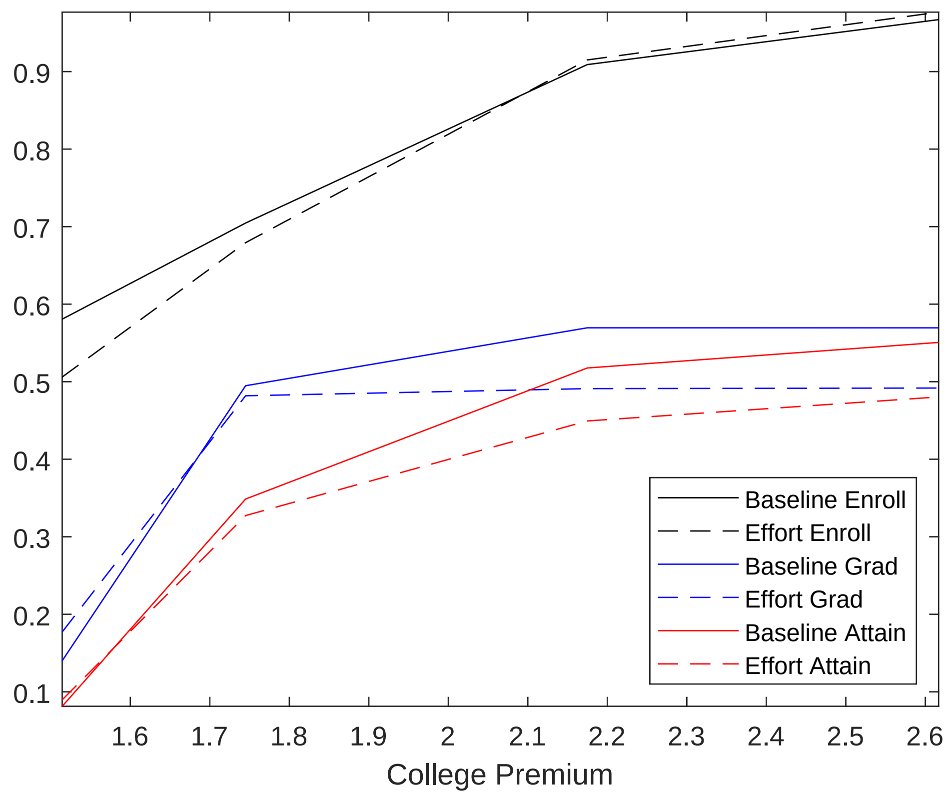

Figure 14: Attainment and Enrollment Across Premia: Benchmark Model vs. Effort Choice

The results are plotted in 14. The punchline is clear: The model that allows for effort choice produces very similar enrollment and attainment levels across skill premia as that which does not. In one sense, this is perhaps to be expected. The reason is that even under a low college premium, the benefits of completing college in the model are substantial and the opportunity costs of attendance are high. Thus, barring extraordinary disutility of ensuring the best odds of completion given one's preparedness, it is natural to expect effort from enrollees. Given that this richer model does not change the message of the paper in a quantitatively important way, we retain the simpler model as our baseline.

Lastly, note that the premium that agents expect is the one that they individually would realize if they completed, which selection bias may prevent from coinciding with the observed premium. For one thing, those with high productivity may decide to leave college, or not enroll at all, but will as a result have high earnings as unskilled households, which would depress the true college premium, all else equal. Hendricks and Leukhina (2012) point out another source of selection bias. Failure risk, when (negatively) correlated with future productivity, will contribute additional bias for two reasons. First, among those who enroll in college, high-ability students graduate at higher rates and subsequently earn higher payoffs than those who failed would have. Second, ex ante, higher-ability individuals, because they face higher probabilities of completion, will, all else equal, enroll at higher rates than lower-ability individuals. Thus, the pool of potential college grades is already selected in favor of high ability. Hendricks and Leukhina (2011) argue that "risk selection" accounts for up to half of the observed premium. This bias does not alter our basic conclusions, primarily because calibrating to match the lower premium implied by their measure still requires matching observed enrollment rates, which are high enough to leave most insenitive to changes in the premium.